

Cisco Security Cloud Control: Secure Firewall Threat Defense Management for Government

First Published: 2022-06-28

Last Modified: 2026-09-18

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2022–2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

PART I **Cisco Security Cloud Control Management: Cloud-Delivered Firewall Management Center** 95

CHAPTER 1 **Configure Cloud-Delivered Firewall Management Center-Managed Secure Firewall Threat Defense** 1

Overview of Cloud-Delivered Firewall Management Center 1

Navigate to the Cloud-Delivered Firewall Management Center in your Security Cloud Control Tenant 3

Determine the Cloud-Delivered Firewall Management Center Version in Security Cloud Control Firewall Management 3

Enable Cloud-Delivered Firewall Management Center 4

Hardware and Software Support 5

Security Cloud Control Firewall Management Platform Maintenance Schedule 5

Licensing 6

Integrate an On-Premises Firewall Management Center With Cisco Security Cloud 6

Security Cloud Control Firewall Management Integrations Page 9

CHAPTER 2 **Manage Multicloud Defense-Onboarded Secure Firewall Threat Defense Virtual Devices** 13

Overview of Multicloud Defense-Onboarded Firewall Threat Defense Virtual Devices 13

Onboard and Configure a Secure Firewall Threat Defense Virtual Device in Multicloud Defense 15

PART II **Onboard Devices to Cloud-Delivered Firewall Management Center** 17

CHAPTER 3 **Onboard a Secure Firewall Threat Defense to the Cloud-Delivered Firewall Management Center** 19

Onboarding Overview 19

Maximum Firewall Threat Defense devices supported in Cloud-Delivered Firewall Management Center 20

Prerequisites to Onboard a Device to Cloud-Delivered Firewall Management Center 21

Onboard a Device with a CLI Registration Key 22

Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning 24

Onboard a Firewall Threat Defense Device to On-Prem Firewall Management Center using Zero-Touch Provisioning 25

Onboard Firewall Threat Defense Devices using Device Templates to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning 27

Cloud Deployment 29

Configure FTDv enforcement point 29

Configure Multicloud Defense enforcement point 37

Onboard a Secure Firewall Threat Defense Cluster 46

Onboard a Chassis 49

Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI 51

Delete Devices from Cloud-Delivered Firewall Management Center 54

Troubleshooting 55

Troubleshoot Cloud-Delivered Firewall Management Center Connectivity with TCP 55

Troubleshoot Firewall Threat Defense Device Connectivity 56

Troubleshoot Device Connectivity Loss After Cloud-Delivered Firewall Management Center Update 56

Troubleshoot Onboarding a Device to the Cloud-Delivered Firewall Management Center Using the CLI Registration Key 57

Error: Device Remains in Pending Setup State After Onboarding 57

Troubleshoot Onboarding a Device to Cloud-Delivered Firewall Management Center Using the Serial Number 58

Device is Offline or Unreachable 58

Error: Serial Number Already Claimed 58

Error: Claim Error 60

Error: Failed to Claim 60

Error: Provisional Error 61

CHAPTER 4 **Migrate On-Premises Firewall Management Center Managed Secure Firewall Threat Defense to Cloud-delivered Firewall Management Center 63**

Overview of Firewall Threat Defense to Cloud-Delivered Firewall Management Center migration 63

How Firewall Threat Defense Migration to Cloud-Delivered Firewall Management Center Works 64

About 14-day evaluation period 65

How Firewall Threat Defense licenses are handled during migration to Cloud-Delivered Firewall Management Center 65
End-to-End Workflow for Migrating Firewall Threat Defense to Cloud-Delivered Firewall Management Center 66
Event and Analytics Management options during Firewall Threat Defense migration 67
Supported features 68
Unsupported features 70
User Identity Migration Guidelines and Limitations for Firewall Threat Defense Devices 71
Migration guidelines and limitations for VPN configuration 71
Supported On-Premises Firewall Management Center and Firewall Threat Defense software versions for migration 73
Prerequisites for Migrating Firewall Threat Defense to Cloud-Delivered Firewall Management Center 75
Migrate Firewall Threat Defense to Cloud-Delivered Firewall Management Center 78
View a Firewall Threat Defense Migration Job 81
Proceed Migration Process 82
Commit Migration Changes Manually to Cloud-Delivered Firewall Management Center 83
Revert the Firewall Threat Defense Management to On-Premises Firewall Management Center 84
View Migrated Devices 84
Generate a Firewall Threat Defense Migration Report 86
Delete a Migration Job 86
Enable Notification Settings 86
Verify Firewall Threat Defense Connectivity with Cloud-Delivered Firewall Management Center 87
Use the Troubleshoot Feature for Failed Migrations 88
Resolving Firewall Threat Defense Migration to Cloud-Delivered Firewall Management Center Issues 90

CHAPTER 5 **Device Registration** 99
Log into the Command-Line Interface on the device 99
Device registration management 101
About the device management page 101
Add a Device Group 102
Device management page enhancements 103
Use the enhanced filters in device management page 105
Add a device group 105

Move devices to another group 106
Managed device list download 107
Shut down or restart the device 107
Download the managed device list 108
Migrate Firewall Threat Defense devices 108
Supported Devices for Migration 109
License for Firewall Threat Defense migration 110
Prerequisites for migration 110
Configurations that the wizard migrates 112
Guidelines and limitations for migration 113
Migrate a secure Firewall Threat Defense 115
Best practices for Threat Defense device migration 117
Use the Security Cloud Control command line interface tool 118

CHAPTER 6 **Device Registration Using Device Templates** 119
Device registration using device templates 119
Variables and network object overrides 119
Model mapping 120
Templates and high availability 120
Validating template configurations before and after application on devices 121
Prerequisites for device registration using device templates 121
Licenses for device registration using device templates 122
Guidelines for device registration using device templates 123
Configure device templates 126
Add a device template 126
Create a new device template 126
Generate a new device template from an existing device 126
Import a device template 127
Configure device settings in the template 128
Add a physical interface 128
Add a logical interface 128
Edit an interface 129
Configure other device settings 130
Template settings 130

Edit general settings 131
Edit licenses 131
Edit applied policies 132
Edit advanced settings 132
Edit deployment settings 133
Configure template parameters 134
Add model mapping 137
Configure Site-to-Site VPN connections in a device template 138
Configure an SD-WAN VPN connection 138
Configure a route-based Site-to-Site VPN connection 139
Configure a policy-based Site-to-Site VPN connection 140
Add a device to an SD-WAN topology in a dual ISP deployment 142
Configure a Template for Firewall Threat Defense Devices Managed Using the Data Interface 143
Use templates with devices 145
Apply templates to existing devices 145
Apply a template 145
Reapply a template 146
Monitor device templates 147
View associated devices 147
Template Apply reports 148
Audit logs 148
Troubleshoot device templates 149

CHAPTER 7 **Transparent or Routed Firewall Mode 153**
Firewall mode 153
Routed firewall mode 153
Transparent firewall mode 154
Transparent firewall network usage 154
Passing traffic for routed-mode features 155
Bridge groups 155
Bridge virtual interfaces 155
Bridge groups in transparent firewall mode 156
Bridge groups in routed firewall mode 157
Allow layer 3 traffic 158

Allowed MAC addresses 158
BPDUs handling 158
MAC address and route lookups 159
Unsupported features for bridge groups in transparent mode 161
Unsupported features for bridge groups in routed mode 162
Default settings 163
Guidelines for firewall mode 163
Set the firewall mode 164

CHAPTER 8 Change Management 167
Change management 167
Make configuration changes in the change management workflow 167
Create separate approver and configuration roles 168
Policies and objects that support change management 169
Requirements and prerequisites for change management 171
Guidelines and limitations for change management 172
Enabling or disabling change management 172
Manage tickets 173
Create change management tickets 175
Opening a ticket for configuration changes 176
Preview a ticket 176
Submit a ticket 177
Discard a ticket 177
Approving or rejecting a ticket 178
Take over or reassign tickets 179

CHAPTER 9 CLI Users 181
CLI users 181
Internal and external users 181
CLI access 182
CLI user roles 182
Guidelines for CLI users 183
Add an internal user at the CLI 183
Troubleshoot LDAP authentication connections 186

CHAPTER 10 **Configuration Deployment** 189

Configuration deployment 189

Configuration changes that require deployment 189

Deployment preview 190

Guidelines for selective policy deployment 191

System username 193

Auto-enabling of application detectors 194

Asset rediscovery with network discovery policy changes 194

Snort restart scenarios 194

Inspection interruption warnings during deployment 195

Inspect traffic during policy apply 196

Snort restart traffic behavior 196

Configurations that restart the snort process when deployed or activated 198

Actions that immediately restart the snort process 199

Requirements and prerequisites for deployment management 200

Best practices for deploying configuration changes 200

Configuration deployment 201

Deploy configuration changes 202

Redeploy existing configurations to a device 208

Deployment management 209

Deployment status 209

View deployment history 210

Deployment rollback 214

Perform a rollback 216

View the deployment rollback transcript 217

Download policy changes report for multiple devices 218

Compare policies 219

Generate current policy reports 220

PART III **System Settings** 223

CHAPTER 11 **System Configuration** 225

Meet requirements and prerequisites for the system configuration 225

Manage the Cloud-Delivered Firewall Management Center system configuration 226
Configure access control preferences 226
Audit logs 227
Stream audit logs to syslog 228
Stream audit logs to an HTTP server 230
Audit log certificates 231
Secure audit log streaming 232
Obtain a signed audit log client certificate for the Cloud-Delivered Firewall Management Center 233
Import an audit log client certificate into the Cloud-Delivered Firewall Management Center 234
Require valid audit log server certificates 235
View the audit log client certificate on the Cloud-Delivered Firewall Management Center 237
Change reconciliation 237
Configure change reconciliation 238
Change reconciliation options 239
Email notifications 239
Intrusion policy preferences 239
Set intrusion policy preferences 240
Network analysis policy preferences 241

CHAPTER 12 Users 243
CLI users 243
Internal and external users 243
User roles 244
Create a User Record with Your Security Cloud Control Username 246
External authentication configuration for Cloud-Delivered Firewall Management Center 247
External authentication objects for Cloud-Delivered Firewall Management Center 247
LDAP 248
RADIUS 248
Add an LDAP external authentication object for the Security Cloud Control 249
Add a RADIUS external authentication object for Security Cloud Control 255
Troubleshoot LDAP authentication connections 261
Configure User Preferences 261
Change the web interface appearance 261
Set your default time zone 262

CHAPTER 13 **Updates** 263

Content updates 263

Product upgrades 264

Guidelines for content updates 264

Vulnerability database (VDB) updates 264

Manually update the VDB 265

Geolocation database (GeoDB) updates 266

Schedule GeoDB updates 267

Manually update the GeoDB 267

Intrusion rule updates 268

Schedule intrusion rule updates 269

Manually update intrusion rules 269

Guidelines for importing custom Snort 2 rules 270

Import custom Snort 2 rules 271

Intrusion rule update logs 272

CHAPTER 14 **Licenses** 275

Licenses 275

Smart Software Manager and accounts 276

Licensing for Management Center and devices 276

Periodic communications with the Smart Software Manager 276

Cloud-Delivered Firewall Management Center licensing and registration 276

Out-of-compliance state 279

Unregistered state 279

End-user license agreement 279

License types and restrictions 279

Essentials licenses 281

Malware Defense Licenses 282

IPS licenses 282

Carrier license 283

URL Filtering licenses 284

Secure Client licenses 285

License for export-controlled functionality 286

Firewall Threat Defense Virtual Licenses 287
Prerequisites for licensing 289
Prerequisites for licensing in high availability, clustering, and multi-instance 289
License requirements for Firewall Threat Defense device high availability 289
License for device clusters 289
Create a Cisco account 290
Create a smart account and add licenses 291
Configure Smart Licensing 292
Register the Cloud-Delivered Firewall Management Center for Smart Licensing 292
Register the Cloud-Delivered Firewall Management Center with the Smart Software Manager 293
Assign licenses to devices 295
Assign licenses to a single device 295
Assign licenses to multiple managed devices 296
Manage Smart Licensing 297
Deregister the Cloud-Delivered Firewall Management Center 297
Monitor Smart License status 298
Monitor Smart Licensing 298
Troubleshoot Smart Licensing 299
Configure Legacy Cloud-Delivered Firewall Management Center PAK-Based Licenses 300
Additional information about licensing 301
Overview 301
Requirements 302
Supported Platforms and Versions 302
How to Implement Universal Permanent Licensing 302
Prerequisites 303
Verify that your Smart Account is Ready to Deploy Universal Permanent License Reservation 304
Enable the Permanent Licensing Menu Option 305
Generate a Reservation Request Code from the Firepower Management Center 306
Generate a Reservation Authorization Code from Cisco Smart Software Manager 306
Enter the Reservation Authorization Code into the Firepower Management Center 307
Apply Permanent Licenses to Managed Firepower Threat Devices 307
Maintain your air-gapped deployment 308
Universal Permanent License Status 308
Maintain Your UPLR Deployment 309

Deactivate and Return a Universal Permanent License 309
Troubleshooting 312
For Assistance 313

CHAPTER 15 **Security certifications compliance 315**
Security certifications compliance modes 315
Security certifications compliance characteristics 316
Security certifications compliance recommendations 317
Appliance hardening 318
Protect your network 319

PART IV **Optimize Firewall Performance with AgenticOps 321**

CHAPTER 16 **Introduction to AgenticOps insights 323**
About AgenticOps insights 323
AgenticOps licensing requirements 325
Prerequisites to use AgenticOps 325
View summary dashboard 325
View AgenticOps insights 329
AgenticOps for On-Premises Firewall Management Center 330
About capacity and traffic insights 331
Analyze capacity trends and forecast resource usage 333
About policy change impact insights 334
Analyze policy change impact 338
About best practices and recommendations 339
Implement best practices and recommendations 340
About feature adoption insights 342
Assess and improve feature adoption 342
About software upgrade insights 343
Manage software upgrades 344
About application insights 345
Detect application outages 346
About compliance insights 346
Generate compliance reports 347

Manage insight preferences 349
Enable AgenticOps insights 349
Disable AgenticOps insights 349
Manage AgenticOps data sources 350
Capacity and traffic insights 351
Best practices and recommendations insights 352
Policy change impact analyzer insights 352
Policy change impact analyzer insights 353
Operational insights 353
Application insights 354
Frequently asked questions about AgenticOps 354
Additional resources 355

PART V **Simplify Firewall Operations with Agent Workforce** 357

CHAPTER 17 **Introduction to Agent Workforce** 359
About Agent Workforce 359
Available agents in Agent Workforce 359
Manage Conversations in Agent Workforce 360
Limitations and considerations 361
Troubleshoot site-to-site VPN issues with VPN agent 361
Benefits of using the VPN agent 361
How VPN agent works 361
Best practices for using the VPN agent 362
Example prompts and expected outcomes for the VPN agent 363
Remediate traffic congestion with Elephant flow agent 363
Benefits of using the Elephant flow agent 364
How Elephant flow remediation works 364
Best practices for using the Elephant flow agent 365
Analyze firewall data with Firewall insight agent 365
Benefits of using the Firewall insight agent 365
How Firewall insight agent works 365
Best practices for using the Firewall insight agent 366
Example prompts and expected outcomes for the Firewall insight agent 366

PART VI Health and Monitoring 369

CHAPTER 18 Health 371

Requirements and prerequisites for health monitoring 371
Health monitoring in Cloud-Delivered Firewall Management Center 371
Health modules 373
Configure health monitoring 383
Health policies 383
Default health policy 384
Create a health policy 385
Apply a health policy 385
Edit a health policy 386
Set a default health policy 388
Delete a health policy 388
Device exclusion in health monitoring 389
Exclude devices from health monitoring 390
Exclude health policy modules on a device 391
Exclude interface health monitoring 392
Manage expired health monitor exclusions 392
Health monitor alerts 393
Health monitor alert information 393
Health alerts for Secure Firewall 200 Series device 394
Health event external notifications 395
Create a health monitor alert 396
Edit a health monitor alert 397
Delete health monitor alerts 397
About the health monitor 398
Device health monitors 399
View system details and troubleshoot devices 400
View the device health monitor 401
Cluster health monitor 404
View the cluster health monitor 405
Health monitor status categories 407

Health event views 407
View health events 408
View the health events table 408
Health event fields 409
Sending device health metrics to OpenConfig 410
Generate certificates and private keys 411
Configure OpenConfig streaming telemetry 413
Troubleshoot OpenConfig streaming telemetry 414
system auditing 416
Audit records 416
Audit log workflow fields 416
The audit events table view 417

CHAPTER 19 Troubleshooting 419
Guidelines and best practices for troubleshooting 419
System messages 420
Message types 420
Message management 422
View basic system information 422
View appliance information 423
Manage system messages 423
View deployment messages 424
View upgrade messages 425
View health messages 425
View task messages 426
Manage task messages 427
Memory usage thresholds for health monitor alerts 427
Disk usage and drain of events health monitor alerts 428
Clear disk space 432
Health monitor reports for troubleshooting 433
Generate troubleshooting files for specific system functions 433
Download advanced troubleshooting files 434
Enhanced troubleshooting experience using Cisco RADKit integration 435
Guidelines and limitations for RADKit integration 435

Enroll RADKit service 436
Manage RADKit service authorization 436
Enable sudo access for devices 438
Download session logs 438
General troubleshooting 439
Connection-based troubleshooting 439
Troubleshoot a connection 439
Advanced troubleshooting for the Secure Firewall Threat Defense device 440
Packet capture 440
Use a packet capture trace 442
Packet tracer tool overview 444
Run a packet trace 444
CPU profiler overview 448
Use the CPU profiler 448
Rule profiler overview 450
Use the rule profiler 450
Use the Firewall Threat Defense diagnostic CLI from the web interface 451
Feature-specific troubleshooting 453

PART VII **Tools** 455

CHAPTER 20 **Backup/Restore** 457
About Backup and Restore 457
Requirements for Backup and Restore 458
Guidelines and Limitations for Backup and Restore 459
Best Practices for Backup and Restore 460
Back Up Managed Devices 462
Back Up a Firewall Threat Defense Device from Cloud-delivered Firewall Management Center 463
Restore Security Cloud Control-Managed Devices 464
Restore a Firewall Threat Defense Device 464
Restore Firewall Threat Defense from Backup: Firewall Threat Defense Virtual 467

CHAPTER 21 **Scheduling** 471
Task scheduling 471

Guidelines for scheduled tasks 471
Scheduled task types 472
Schedule a task 474

CHAPTER 22 Import/Export 477
Import/Export feature overview 477
Exceptions to export behavior 477
Import objects and object groups 478
Conflict resolution for duplicate configurations 479
Requirements and prerequisites for import/export 480
Guidelines for Import/Export 480
Export configurations 481
Import configurations 482

PART VIII Reporting and Alerting 485

CHAPTER 23 External Alerting with Alert Responses 487
Configuring external alerts with alert responses 487
External alert types 488
Prerequisites for external alerting with alert responses 488
Guidelines for notifying external systems with alert responses 488
Alert responses 489
Create an SNMP alert response 489
Syslog alert responses 491
SYSLOG alert facilities 491
Syslog severity levels 492
Create a syslog alert response 493
Create an email alert response 493
Create a webhook alert response 494
Troubleshoot external alerts and alert response configuration 495

CHAPTER 24 External Alerting for Intrusion Events 497
External alerts for intrusion events 497
License requirements for external alerting for intrusion events 498

Requirements for configuring external alerting for intrusion events 498
Configure SNMP alerting for intrusion events 498
Intrusion SNMP alert configuration options 499
Configure syslog alerting for intrusion events 500
Facility and severity values for intrusion syslog alerts 501
Configure email alerting for intrusion events 503
Intrusion email alert options 503

PART IX Event and Asset Analysis Tools 505

CHAPTER 25 Unified Events 507
Unified events 507
Work with unified events 507
Set a time range in unified events 510
Filters in unified events 511
Save a search in unified events 512
Load a saved search in unified events 513
Save a column set 513
Load a saved column set 514
Unified events column details 514

CHAPTER 26 Lookups 517
Lookups in Cloud-Delivered Firewall Management Center 517
Lookup IP address details using Whois 517

CHAPTER 27 Event searches 519
Search constraints 520
Best practice for general search constraints 520
Wildcards and symbols in searches 521
Objects and application filters in searches 521
Time constraints in searches 521
IP addresses in searches 522
URLs in searches 523
Managed devices in searches 523

Ports in searches 523
Perform a search 524
Save a search 524
Load a saved search 524

CHAPTER 28 Send Firewall Threat Defense Events to Cisco Splunk 527
Splunk integration: Send events directly from Cloud-Delivered Firewall Management Center 527
Guidelines and limitations for Splunk integration 527
Configure Splunk in Cloud-Delivered Firewall Management Center 528
Configure Splunk server 530
Select event types 530
Select devices and interfaces 531
Configure Firewall certificates for TLS 532
Review and submit the Splunk profile 532
Configure Secure Firewall App in Splunk 533

CHAPTER 29 Event Investigation Using Web-Based Resources 535
Event investigation using web-based resources 535
Manage contextual cross-launch resources 536
Requirements for custom contextual cross-launch resources 536
Add contextual cross-launch resources 536
Investigate events using contextual Cross-Launch 537

PART X Events and Assets 539

CHAPTER 30 Monitor Cloud-Delivered Firewall Management Center-Managed Threat Defense Device Events 541
About Security Analytics and Logging 541
Comparison of SAL Remote Event Storage and Monitoring Options 542
About SAL (OnPrem) 542
Licensing for SAL (OnPrem) 543
Manage SAL (OnPrem) for Security Cloud Control-Managed Firewall Threat Defense Devices 543
Configure SAL (OnPrem) Integration 545
Configure a Secure Network Analytics Manager 545
Configure a Secure Network Analytics Data Store 546

About Security Analytics and Logging(SaaS) for Firewall Threat Defense 548
Licensing for SAL (SaaS) 548
Configure the SAL (SaaS) Integration 548
Requirements and Guidelines 549
Send Cloud-Delivered Firewall Management Center-Managed Events to SAL (SaaS) Using Syslog 553
Send Cloud-Delivered Firewall Management Center-Managed Event Logs to SAL (SaaS) Using a Direct Connection 555
Enable Individual Threat Defense Devices to Send Events to SAL (SaaS) Using a Direct Connection 556
View and Work with Events in Security Cloud Control 557
View AI Defense Events in Security Cloud Control 558

CHAPTER 31 FTD Dashboard 559
About the FTD Dashboard 559
View the FTD Dashboard 560
FTD Dashboard Widgets 561
Top Intrusion Rules Widget 561
Top Intrusion Attackers Widget 561
Top Intrusion Targets Widget 562
Top Malware Signatures Widget 562
Top Malware Senders Widget 562
Top Malware Receivers Widget 562
Malware Events by Disposition Widget 562
Network Activity Widget 562
Event Activity Widget 562
Access Control Actions Widget 562
Top Access Control Policies Widget 562
Top Access Control Rules Widget 563
Top Devices Widget 563
Top Users Widget 563
Top Users by Blocked Connections Widget 563
Top Devices with Health Alerts Widget 563
Top Loaded Devices Widget 563

Top Web Applications Widget 563
Top Client Applications Widget 563
Top Blocked Web Applications Widget 563
Modify Time Settings for the FTD Dashboard 564

CHAPTER 32 **Connection Logging** 565
Connection logging 565
Logged connections 566
Other Connections You Can Log 567
Rules and policy actions effect on logging 568
Logging for fastpathed connections 568
Logging for monitored connections 568
Logging for trusted connections 569
Logging for blocked connections 569
Logging for allowed connections 571
Connection logging modes 572
Connection logging limitations 574
Best practices for connection logging 574
Requirements and prerequisites for connection logging 576
Connection logging 576
Logging connections with tunnel and prefilter rules 576
Configure logging for decryptable connections with TLS/SSLDecryption rules 577
Logging connections with security intelligence 578
Logging connections with access control rules 579
Logging connections with a policy default action 580
Limiting logging of long URLs 580
Application-aware and protocol-aware syslogs 581
Advanced logging guidelines and limitations 582
Enable advanced logging 582
Configure protocols for advanced logging 584
Monitor and troubleshoot advanced logging 585

CHAPTER 33 **Advanced Logging Syslog Fields** 587
Common fields 587

CONN protocol fields 588
DNS protocol fields 589
FTP fields 590
HTTP fields 591
Notice protocol fields 592
Weird protocol fields 592

CHAPTER 34 User Identity and Activity 595
User-related fields 595
Active sessions data 602
View active session data 602
User data 603
View user data 605
User activity data 606
View user activity data 608
User profile and host history 608
View user details and host history 609

PART XI High Availability and Scalability 611

CHAPTER 35 Multi-Instance Mode 613
Multi-instance mode 613
Multi-instance modes and appliance modes 613
Chassis management interface 614
Instance interfaces 614
Interface types 615
Chassis interfaces vs. instance interfaces 615
Shared interface scalability 617
Shared interface best practices 618
How the chassis classifies packets 620
Classification examples 621
Cascading container instances 621
Typical multi-instance deployments 622
Automatic MAC addresses for instance interfaces 623

Performance scaling factors for multi-instance modes 624
Instances and high availability 624
Licenses for instances 625
Requirements and prerequisites for instances 625
Guidelines and limitations for instances 626
Instance configuration 629
Convert a device to Multi-Instance mode 629
Chassis interface configuration 633
Configure a physical interface 633
Configure an EtherChannel 636
Configure a subinterface 640
Add an instance 642
System configuration customization 649
Configure SNMP 649
Import or export the chassis configuration 650
Chassis platform settings 652
Create a chassis platform settings policy 652
Configure DNS 653
Configure SSH and SSH access list 654
Configure syslog 659
Configure time synchronization 664
Configure time zones 666
Multi-instance mode management 667
Onboard the multi-instance chassis using the CLI 667
Change interfaces assigned to an instance 671
Change chassis management settings at the FXOS CLI 673
Multi-instance mode 676
Multi-instance setups 676
Instance interfaces 677

CHAPTER 36 **High Availability for Devices** 681
Secure Firewall Threat Defense high availability 681
High availability support on Firewall Threat Defense devices in a remote branch office deployment 682

High availability system requirements 682
Hardware requirements 682
Software requirements 683
License requirements for Firewall Threat Defense devices in a high availability pair 684
Failover and stateful failover links 684
Failover link 684
Stateful failover links 686
Failover and data link interruption avoidance 687
MAC addresses and IP addresses in High availability 688
Stateful failover 690
Supported features 691
Unsupported features 693
Bridge group requirements for high availability 693
Failover health monitoring 694
Unit health monitoring 694
Heartbeat module redundancy 695
Interface monitoring 695
Failover triggers and detection timing 697
Active/standby failover 698
Primary/secondary roles and active/standby status 699
Active unit determination at startup 699
Failover events 699
Config-sync optimization 700
Requirements and prerequisites for high availability 702
Guidelines for High availability 702
Add a high availability pair 705
Configure optional high availability parameters 708
Configure standby IP addresses and interface monitoring 708
Edit high availability failover criteria 709
Configure virtual MAC addresses 709
Manage High availability 710
Switch the Active Peer in the Firewall Threat Defense High Availability Pair 710
Refresh node status for a single Firewall Threat Defense high availability pair 711
Suspend and resume high availability 711

Replace a unit in Firewall Threat Defense high availability pair 712
Replace a Primary Firewall Threat Defense HA Unit with no Backup 712
Replace a Secondary Firewall Threat Defense HA Unit with no Backup 713
Break a high availability pair 714
Remove a High Availability Pair 716
Monitoring High availability 717
View failover history 717
View stateful failover statistics 717
High availability break troubleshooting in remote branch deployments 718
Break a high availability pair in active-active state 718
Break a high availability pair when active or standby unit has lost connectivity 720
Break a high availability pair when the secondary device is in a failed or disabled state 721

CHAPTER 37 **Clustering: Secure Firewall 3100/4200** 723
Clustering for the Secure Firewall 3100/4200 723
How a cluster fits into your network 723
Control and data node roles 724
Cluster interfaces 724
Cluster control links 724
Configuration replication 724
Management network 725
Licenses for clustering 725
Prerequisites for clustering 725
Guidelines for clustering 726
Clustering configuration 731
Cluster interfaces 731
Cluster control link 731
Spanned EtherChannels (Recommended) 734
Individual interfaces 737
Cable and add devices to the Cloud-Delivered Firewall Management Center 739
Create a Cluster 741
Configure interfaces 748
Configure spanned EtherChannels 748
Configure individual interfaces 749

Configure cluster health monitor settings 751
Configure distributed Site-to-Site VPN 756
Distributed site-to-site VPN 756
License for distributed site-to-site VPN 758
Prerequisites for distributed site-to-site VPN 758
Guidelines for distributed site-to-site VPN 759
Enable distributed site-to-site VPN 760
Redistribute distributed site-to-site VPN sessions 760
Cluster node management 762
Add a new cluster node 762
Break a node 764
Break the cluster 765
Disable clustering 766
Rejoin the cluster 767
Change the control node 767
Edit the cluster configuration 768
Reconcile cluster nodes 769
Unregister the Cluster or Nodes and Register to a New Cloud-Delivered Firewall Management Center 771
Monitor the cluster 772
Cluster health monitor dashboard 774
View cluster health 775
Cluster metrics 777
Troubleshoot the cluster 777
Perform a ping on the cluster control link 778
Clustering examples 780
Firewall on a stick 780
Traffic segregation 781
Reference for Clustering 782
Firewall Threat Defense features and clustering 782
Unsupported features with clustering 782
Centralized features for clustering 783
Connection settings and clustering 784
FTP and Clustering 784

Multicast routing in individual interface mode 784
NAT and clustering 784
Dynamic routing 786
Dynamic routing in individual interface mode 787
SIP inspection and clustering 788
SNMP and clustering 788
Syslog and clustering 788
Cisco TrustSec and clustering 788
VPN and clustering 788
Performance scaling factor 789
Control node election 789
High availability within clusters 790
Node health monitoring 790
Interface monitoring 790
Status after failure 791
Rejoin the cluster 791
Data path connection state replication 791
How the cluster manages connections 792
Connection roles 792
New connection ownership 793
Sample data flow for TCP 794
Sample data flows for ICMP and UDP 795
History for Clustering 796

CHAPTER 38 **Clustering: Private Cloud** 799
About clustering: private cloud 799
How the cluster fits into your network 799
Control and data node roles 800
Individual interfaces 800
Policy-based routing 801
Equal-cost multi-path routing 801
Cluster Control Link 802
Cluster control link traffic overview 802
Configuration replication 803

Management network 803
Licenses for clustering: Private cloud 803
Prerequisites for clustering: Private cloud 804
Guidelines for clustering: Private cloud 805
Configure clustering: private cloud 806
Add nodes to the Management Center 806
Create a cluster 806
Configure interfaces 813
Configure Cluster Health Monitor Settings 814
Manage cluster nodes 817
Add a new cluster node 817
Break a node 819
Break the cluster 820
Disable clustering 821
Rejoin the cluster 822
Change the control node 822
Edit the cluster configuration 823
Reconcile cluster nodes 824
Delete the cluster or nodes from the Management Center 825
Monitor cluster status 826
Cluster health monitor 828
View the cluster health monitor 829
Reference for clustering 832
Threat Defense features and clustering 832
Unsupported features and clustering 832
Centralized features for clustering 832
Connection settings and clustering 833
Dynamic routing and clustering 833
FTP and clustering 834
NAT and clustering 834
SIP inspection and clustering 836
SNMP and clustering 836
Syslog and clustering 836
Cisco trustsec and clustering 836

VPN and clustering 836
Performance scaling factor 837
Control node election 837
High availability within the cluster 838
Node health monitoring 838
Interface monitoring 838
Status after failure 838
Rejoining the cluster 839
Data path connection state replication 839
How the cluster manages connections 840
Connection roles 840
New connection ownership 841
Sample data flow for TCP 841
Sample data flow for ICMP and UDP 842
History for clustering: private cloud 844

CHAPTER 39 **Clustering: Public Cloud** 845
About Threat Defense Virtual clustering in the public cloud 845
How the cluster fits into your network 846
Individual interfaces 846
Control and data node roles 847
Cluster Control Link 847
Cluster control link traffic overview 848
Configuration replication 848
Management network 848
Licenses for Threat Defense Virtual clustering 848
Requirements and prerequisites for Threat Defense Virtual clustering 849
Guidelines for Threat Defense Virtual clustering 851
Deploy the cluster in AWS 852
AWS gateway load balancer and Geneve single-arm proxy 852
Sample topology 853
End-to-end process for deploying Threat Defense Virtual cluster on AWS 854
Templates 856
Deploy the Stack in AWS Using a CloudFormation template 856

Manual cluster deployment in AWS 869
Create the Day0 configuration for AWS 869
Deploy cluster nodes 871
Deploy the cluster in Azure 872
Sample topology for GWLB-based cluster deployment 872
Azure gateway load balancer and paired proxy 873
End-to-end process for deploying Threat Defense Virtual cluster in Azure with GWLB 874
Templates 876
Prerequisites for Azure GWLB cluster deployment 876
Deploy cluster on Azure with GWLB using an Azure resource manager template 876
Sample topology for NLB-based cluster deployment 879
End-to-end process for deploying Threat Defense Virtual cluster in Azure with NLB 880
Templates 882
Prerequisites for Azure NLB cluster deployment 882
Deploy cluster on Azure with NLB using an Azure resource manager template 883
Manual cluster deployment in Azure 885
Create the Day0 configuration for Azure 885
Deploy cluster nodes manually - GWLB-based deployment 887
Deploy cluster nodes manually - NLB-based deployment 888
Troubleshoot cluster deployment in Azure 889
Deploy the cluster in GCP 891
Sample topology for GCP clustering autoscale solution 891
End-to-end process for deploying Threat Defense Virtual cluster in GCP 892
Templates 893
Deploy the instance group in GCP using an instance template 893
Manual cluster deployment in GCP 895
Create the Day0 configuration with a fixed configuration for GCP 895
Deploy cluster nodes manually 896
Allow health checks for GCP network load balancers 897
Add the cluster to the Management Center (manual deployment) 898
Onboard a Secure Firewall Threat Defense Cluster 904
Edit cluster health monitor settings 907
Manage cluster nodes 910
Disable clustering 910

Rejoin the cluster 911
Reconcile Cluster Nodes 911
Unregister the cluster or nodes and register with a new Cloud-Delivered Firewall Management Center 912
Monitor cluster status 913
Cluster health monitor 915
View the cluster health monitor 916
Upgrade the cluster 919
Reference for clustering 919
Threat Defense features and clustering 919
Unsupported features and clustering 920
Centralized features for clustering 920
Cisco Trustsec and clustering 920
Connection settings and clustering 921
Dynamic routing and clustering 921
FTP and clustering 922
NAT and clustering 922
SIP inspection and clustering 924
SNMP and clustering 924
Syslog and clustering 924
VPN and Clustering 924
Performance scaling factor 925
Control node election 925
High availability within the cluster 926
Node health monitoring 926
Interface monitoring 926
Status after failure 926
Rejoining the cluster 927
Data path connection state replication 927
How the cluster manages connections 928
Connection roles 928
New connection ownership 929
Sample data flow for TCP 929
Sample data flow for ICMP and UDP 930

History for Threat Defense Virtual clustering in the public cloud 932

CHAPTER 40 **Clustering: Firepower 4100/9300** 935

About clustering on the Firepower 4100/9300 chassis 935

Bootstrap configuration 936

Cluster members 936

Cluster control link 936

Size the cluster control link 937

Cluster control link redundancy 937

Cluster control link reliability for inter-chassis clustering 938

Cluster control link network 938

Management network 938

Management interface 938

Cluster interfaces 938

Spanned EtherChannels 939

Configuration replication 939

Licenses for clustering 939

Requirements and prerequisites for clustering 940

Clustering guidelines and limitations 944

Configure clustering 947

FXOS: Add a Firewall Threat Defense cluster 947

Create a Firewall Threat Defense cluster 947

Add more cluster nodes 950

Cloud-Delivered Firewall Management Center: Add a cluster 951

Cloud-Delivered Firewall Management Center: Configure cluster, data interfaces 957

Cloud-Delivered Firewall Management Center: Configure Cluster Health Monitor Settings 959

FXOS: Remove a cluster node 962

Cloud-Delivered Firewall Management Center: Manage cluster members 964

Add a new cluster member 964

Replace a cluster member 964

Deactivate a member 966

Rejoin the cluster 967

Unregister a data node 967

Change the control unit 968

Reconcile cluster members 969
Cloud-Delivered Firewall Management Center: Monitor the cluster 969
Cluster Health Monitor Dashboard 972
Viewing cluster health 973
Cluster metrics 975
Examples for Clustering 976
Firewall on a stick 976
Traffic segregation 976
Reference for clustering 977
Firewall Threat Defense features and clustering 977
Unsupported features with clustering 977
Centralized features for clustering 978
Connection settings 978
Dynamic routing and clustering 979
FTP and clustering 979
Multicast routing and clustering 979
NAT and clustering 979
SIP inspection and clustering 981
SNMP and clustering 981
Syslog and clustering 981
TLS/SSL connections and clustering 981
Cisco TrustSec and clustering 981
VPN and clustering 981
Performance scaling factor 982
Control unit election 982
High availability within the cluster 983
Chassis-application monitoring 983
Unit health monitoring 983
Interface monitoring 983
Decorator application monitoring 984
Status after failure 984
Rejoining the cluster 984
Data path connection state replication 985
How the cluster manages connections 985

Connection roles 985
New connection ownership 987
Sample data flow for TCP 987
Sample data flow for ICMP and UDP 988
History for clustering 989

PART XII **Interfaces and Device Settings** 995

CHAPTER 41 **Interface Overview** 997
Management interface 997
Management interface 997
Diagnostic interface 998
Interface Mode and Types 999
Security zones and interface groups 999
Auto-MDI/MDIX feature 1001
Redundant interfaces (deprecated) 1001
Default settings for interfaces 1001
Create security zone and interface group objects 1002
Enable the physical interface and configure Ethernet settings 1003
Configure EtherChannel interfaces 1006
About EtherChannels 1006
EtherChannels 1006
Guidelines for EtherChannels and redundant interfaces 1009
Configure an EtherChannel 1011
Sync Interface Changes with the Cloud-Delivered Firewall Management Center 1014
Network module management for Secure Firewall 3100/4200 1015
Configure breakout ports 1016
Add a network module 1020
Hot swap the network module 1022
Replace the network module with a different type 1024
Remove the network module 1028
Merge the management and diagnostic interfaces 1030
Unmerge the management interface 1034
History for interfaces 1035

CHAPTER 42 **Regular Firewall Interfaces** 1039

Requirements and prerequisites for regular firewall interfaces 1039
Configure switch ports 1040
About switch ports 1040
Switch ports and interfaces 1040
Auto-MDI/MDIX feature 1041
Prerequisites for switch ports 1041
Guidelines for switch ports 1041
Configure switch ports and power over ethernet 1043
Enable or disable switch port mode 1043
Configure a VLAN interface 1044
Configure switch ports as access ports 1046
Configure switch ports as trunk ports 1048
Configure Power over Ethernet 1051
Loopback interfaces 1052
Guidelines for Loopback Interfaces 1053
Configure a loopback interface 1053
Rate-Limit traffic to the loopback interface 1054
VLAN subinterfaces and 802.1Q trunking 1058
Guidelines for VLAN subinterfaces 1058
Maximum number of VLAN subinterfaces by device model 1059
Add a subinterface 1060
Configure VXLAN Interfaces 1062
VXLAN interfaces 1062
encapsulation 1062
VXLAN tunnel endpoints 1062
VTEP source interface 1063
VNI interfaces 1063
How VXLAN and Geneve packets are processed 1064
Peer VTEPs 1065
VXLAN use cases 1066
Requirements and prerequisites for VXLAN interfaces 1070
Guidelines for VXLAN interfaces 1070

Configure VXLAN or Geneve interfaces 1071
Configure VXLAN interfaces 1071
Configure geneve interfaces 1074
Allow gateway load balancer health checks 1076
Configure routed and transparent mode interfaces 1077
Routed and transparent mode interfaces 1077
Dual IP stack (IPv4 and IPv6) 1077
31-bit subnet masks 1077
Guidelines and limitations for routed and transparent mode interfaces 1078
Configure routed mode interfaces 1080
Bridge group interfaces 1086
Configure general bridge group member interface parameters 1086
Configure the bridge virtual interface (BVI) 1088
Configure IPv6 addressing 1090
About IPv6 1090
IPv6 prefix delegation client configuration 1091
Configure a global IPv6 address 1095
Configure IPv6 neighbor discovery 1099
Advanced interface settings 1103
About advanced interface configuration 1103
MAC addresses 1103
MTU 1104
TCP MSS 1106
ARP inspection for bridge group traffic 1107
MAC address tables 1108
Default settings 1108
Guidelines for ARP inspection and the MAC address table 1108
Configure the MTU 1109
Configure the MAC address 1110
Add a static ARP entry 1111
Add a static MAC address and disable MAC learning for a bridge group 1112
Set security configuration parameters 1112
History for Regular Firewall Interfaces 1115

CHAPTER 43 **Inline Sets and Passive Interfaces** 1119

IPS interfaces 1119

Inline sets 1119

Multiple inline pairs and asynchronous routing 1120

Passive interfaces 1121

Hardware Bypass for inline sets 1122

Hardware Bypass triggers 1122

Hardware bypass switchover 1122

Snort Fail Open vs. Hardware Bypass 1123

Hardware Bypass status 1123

Requirements and prerequisites for inline sets 1123

Guidelines for configuring inline sets and passive interfaces 1125

Configure a passive interface 1127

Configure Inline Sets 1128

CHAPTER 44 **Device Settings** 1133

Edit general settings 1133

Copy a configuration to another device 1135

Export and import the device configuration 1136

Edit license settings 1140

View system information 1141

View device inventory 1142

View the inspection engine 1142

Edit health settings 1143

Out-of-band configuration detection 1143

Guidelines for out-of-band configuration 1144

Access recovery-config mode in the diagnostic CLI 1146

Acknowledge the out-of-band configuration 1149

Edit management settings 1153

Configure a redundant manager access data interface 1153

Change manager access interface settings 1158

Change the device IP address 1158

Change the Cloud-Delivered Firewall Management Center IP Address 1169

Change both Cloud-Delivered Firewall Management Center and Threat Defense IP addresses 1171
Change the manager access interface 1174
Troubleshoot the management connection 1181
Manually Roll Back the configuration if the Cloud-Delivered Firewall Management Center loses Connectivity 1181
Troubleshoot Management Connectivity on a Data Interface 1183
Troubleshoot management connectivity on a data interface in a High Availability pair 1184
View inventory details 1186
Edit applied policies 1186
Edit advanced settings 1188
Configure automatic application bypass 1189
Configure object group search 1190
Configure interface object optimization 1192
Edit deployment settings 1192
Edit cluster health monitor settings 1195
Hot swap an SSD 1198
Disable the USB port 1201
Disable the USB port on a device 1201
Disable the USB port in multi-instance mode 1202
Configure SNMP for FXOS 1204
Enable SNMP for FXOS 1204
Create an SNMP Trap for FXOS 1205
Create an SNMP user for FXOS 1207
Configure alarms for the ISA 3000 1207
About alarms 1208
Alarm input interfaces 1208
Alarm output interface 1209
Syslog alarms 1209
SNMP Alarms 1209
Defaults for alarms 1209
Prerequisites for alarms 1210
Configure alarms for the ISA 3000 1210
Configure alarm input contacts 1210
Configure power supply alarms 1213

Configure temperature alarms 1215
Turn off the external alarm 1218
Monitoring alarms 1218
Monitoring alarm status 1218
Monitoring syslog messages for alarms 1218
History for Device Settings 1219

CHAPTER 45 DHCP and DDNS 1223
About DHCP and DDNS services 1223
DHCPv4 servers 1223
DHCP options 1224
DHCPv6 stateless server 1224
DHCP relay agents 1225
Requirements and prerequisites for DHCP and DDNS 1225
Guidelines for DHCP and DDNS services 1225
Configure the DHCPv4 server 1227
DHCPv6 stateless server configuration 1229
Create a DHCP IPv6 pool 1229
Enable the DHCPv6 stateless server 1232
Configure the DHCP relay agent 1233
Configure Dynamic DNS 1234
History for DHCP and DDNS 1239

CHAPTER 46 Platform Settings 1241
Platform settings policies 1242
List requirements and prerequisites for platform settings policies 1242
Manage platform settings policies 1243
Configure ARP inspection 1244
Configure a banner 1245
Configure DNS server settings 1246
Configure external authentication 1249
Configure an LDAP External Authentication Object 1251
Configure a RADIUS Authentication Object 1253
Enable a virtual-router-aware interface for platform external authentication 1255

Configure fragment settings 1256
Configure HTTP access 1257
Configure ICMP access 1258
Configure NetFlow 1260
Add a collector in NetFlow 1260
Add a traffic class to NetFlow 1261
Configure Secure Shell (SSH) access 1262
Configure an SMTP server 1264
Configure SNMP for a security device 1264
SNMP protocol 1266
List SNMP terminology 1266
Describe MIBs and traps 1267
List of supported tables and objects in MIBs 1267
Add SNMPv3 users 1271
Add SNMP hosts 1274
Configure SNMP traps 1276
Configure SSL Settings 1278
Configure SSL settings 1279
Syslog 1283
Syslog messages 1283
Syslog message severity levels 1284
Syslog message filters 1285
Syslog message classes and associated message ID numbers 1285
Guidelines for logging 1288
Configure syslog logging for Firewall Threat Defense devices 1290
Firewall Threat Defense platform settings that apply to security event syslog messages 1291
Configure logging and basic settings 1291
Enable logging destinations 1293
Configure syslog messages to be sent to an e-mail address 1294
Create a custom event list 1295
Configure syslog rate limits 1296
Configure syslog settings 1297
Configure a syslog server 1298
Configure timeouts 1301

Synchronize time on FTD devices 1303
Configure a time zone for a device 1304
Configure UCAPL/CC compliance 1305
Configure a performance profile 1306

CHAPTER 47 FlexConfig policies 1309
FlexConfig policy overview 1309
Recommended usage for FlexConfig policies 1310
CLI commands in FlexConfig objects 1310
Determine the ASA software version and current CLI configuration 1311
Prohibited CLI commands 1312
Template scripts 1314
FlexConfig variables 1314
Variable processing 1315
Evaluate device variable results 1318
FlexConfig policy object variables 1320
FlexConfig system variables 1321
Predefined FlexConfig objects 1322
Predefined text objects 1327
Requirements and prerequisites for FlexConfig policies 1330
Guidelines and limitations for FlexConfig 1331
Customize device configuration with FlexConfig policies 1331
Configure FlexConfig objects 1333
Add a policy object variable to a FlexConfig object 1336
Configure secret keys 1337
Configure FlexConfig text objects 1337
Configure the FlexConfig policy 1339
Set target devices for a FlexConfig policy 1340
Preview the FlexConfig policy 1341
Verify the deployed configuration 1342
Remove features configured using FlexConfig 1344
Convert from FlexConfig to managed feature 1345
Examples for FlexConfig 1346
Configure precision time protocol (ISA 3000) 1346

Configure automatic hardware bypass for power failure (ISA 3000) 1350
Migrate FlexConfig policies 1352

PART XIII **Routing** 1355

CHAPTER 48 **Static and Default Routes** 1357
Static and default routes overview 1357
Default routes 1357
Static routes 1358
Route to null0 interface to drop unwanted traffic 1358
Route priorities 1358
Transparent firewall mode and bridge group routes 1359
Static route tracking 1359
Requirements and prerequisites for static routes 1360
Guidelines for static and default routes 1360
Add a static route 1361
Reference for routing 1362
Path determination 1362
Supported route types 1363
Static versus dynamic routing algorithms 1363
Single-path versus multipath 1363
Flat versus hierarchical routing 1363
Link-state versus distance vector 1364
Supported internet protocols for routing 1364
Routing tables 1365
How the routing table is populated 1365
How forwarding decisions are made 1367
Dynamic routing and High availability 1368
Dynamic routing in Clustering 1368
Dynamic routing in individual interface mode 1369
Routing table for management traffic 1370
Equal-cost multi-path routing 1371
Routing GRE traffic 1372
About route maps 1372

Permit and deny clauses 1372
Match and set clause values 1373

CHAPTER 49 Virtual Routers 1375
Virtual routers and virtual routing and forwarding (VRF) 1375
Virtual router applications 1376
Global and user-defined virtual routers 1376
Virtual-router-aware policies 1377
Virtual router interconnection 1378
Best practice for overlapping IP addresses 1380
Configuring SNMP on user-defined virtual routers 1381
Maximum number of virtual routers by device model 1381
Requirements and prerequisites for virtual routers 1383
Guidelines and limitations for virtual routers 1383
Manage virtual routers 1386
Create a virtual router 1386
Configure a virtual router 1387
Modify a virtual router 1388
Remove Virtual Routers 1389
Monitor virtual routers 1390
Configuration examples for virtual routers 1390
Route to a distant server through virtual routers 1390
How to Provide Internet Access with Overlapping Address Spaces 1395
Allow RA VPN access to internal networks in virtual routing 1402
Secure traffic from networks in multiple virtual routers over a site-to-site VPN 1405
Secure traffic from networks with multiple virtual routers over a site-to-site VPN with dynamic VTI 1409
Route traffic between two overlapping network hosts in virtual routing 1411
Manage overlapping segments in routed firewall mode with BVI interfaces 1414
Configure user authentication with overlapping networks 1418
Interconnect virtual routers using BGP 1424

CHAPTER 50 ECMP 1431
ECMP routing methods 1431

Guidelines and limitations for ECMP 1432
Manage ECMP zones for virtual routers 1433
Create an ECMP zone 1433
Configure an equal cost static route 1434
Modify an ECMP zone 1435
Remove an ECMP zone 1436
Configure ECMP zones on your device 1436

CHAPTER 51 **Bidirectional Forwarding Detection Routing** 1441
BFD routing protocol 1441
Guidelines for BFD routing 1442
Configure BFD 1443
Configure BFD policies 1443
Configure a single-hop BFD policy 1444
Configure a multi-hop BFD policy 1445
History for BFD routing 1446

CHAPTER 52 **OSPF** 1447
OSPF 1447
OSPF routing 1447
OSPF support for fast hello packets 1449
OSPF fast hello packets configuration prerequisites 1449
OSPF hello interval and dead interval 1449
OSPF fast hello packets 1449
Benefits of OSPF fast hello packets 1450
OSPFv2 and OSPFv3 implementation differences 1450
Requirements and prerequisites for OSPF 1451
Guidelines for OSPF 1451
OSPFv2 configuration 1454
Configure OSPF areas, ranges, and virtual links 1454
Configure OSPF redistribution 1457
Configure OSPF Inter-Area filtering 1459
Configure OSPF filter rules 1460
Configure OSPF summary addresses 1461

Configure OSPF interfaces and neighbors 1462
Configure OSPF Advanced Properties 1465
OSPFv3 configuration 1467
Configure OSPFv3 Areas, Route Summaries, and Virtual Links 1467
Configure OSPFv3 redistribution 1470
Configure OSPFv3 summary prefixes 1471
Configure OSPFv3 interfaces, authentication, and neighbors 1472
Configure OSPFv3 Advanced Properties 1475
History for OSPF 1477

CHAPTER 53 EIGRP 1479
EIGRP routing 1479
Requirements and prerequisites for EIGRP 1480
Guidelines and limitations for EIGRP routing 1481
Configure EIGRP routing on firewall devices 1482
Configure EIGRP settings 1483
Configure EIGRP neighbors settings 1483
Configure EIGRP filter rules 1484
Configure EIGRP redistribution settings 1485
Configure EIGRP summary address settings 1486
Configure EIGRP interfaces settings 1487
Configure EIGRP advanced settings 1488

CHAPTER 54 BGP 1491
BGP 1491
Routing tables 1491
BGP usage scenarios 1492
BGP path selection 1493
BGP multipath 1493
Requirements and prerequisites for BGP 1495
Guidelines for BGP 1495
Configure BGP 1496
Configure basic BGP settings 1497
Configure BGP general settings 1499

Configure BGP Neighbor Settings 1500
Configure BGP aggregate address settings 1504
Configure BGPv4 filtering settings 1505
Configure BGP network settings 1506
Configure BGP redistribution settings 1507
Configure BGP route injection settings 1508
Configure BGP route import and export 1509

CHAPTER 55 RIP 1513
RIP 1513
How RIP routing updates work 1514
RIP routing metric 1514
RIP stability features 1514
RIP timers 1515
Requirements and prerequisites for RIP 1516
Guidelines for RIP 1516
Configure RIP 1517

CHAPTER 56 Multicast 1521
Multicast routing 1521
IGMP Protocol 1522
Stub multicast routing 1522
PIM multicast routing 1523
PIM sparse mode 1523
PIM source specific multicast support 1523
Multicast Bidirectional PIM 1524
PIM bootstrap router (BSR) 1524
PIM bootstrap router (BSR) terminology 1525
Multicast group concept 1525
Multicast Addresses 1526
Clustering 1526
Prerequisites for Multicast Routing 1526
Guidelines for multicast routing 1526
Configure IGMP features 1527

Enable multicast routing 1528
Configure IGMP protocol 1529
Configure IGMP access groups 1530
Configure IGMP static groups 1531
Configure IGMP join groups 1531
Configure PIM features 1532
Configure PIM protocol 1533
Configure PIM neighbor filters 1534
Configure PIM bidirectional neighbor filters 1534
Configure PIM rendezvous points 1535
Configure PIM route trees 1536
Configure PIM request filters 1537
Configure the Firewall Threat Defense device as a candidate bootstrap router 1538
Configure multicast routes 1538
Configure multicast boundary filters 1539

CHAPTER 57 **Policy-Based Routing** 1541
About policy-based routing 1541
Licenses for policy-based routing 1543
Best practices for policy-based routing 1543
Path monitoring metrics for best path route determination 1545
Configure path monitoring settings 1547
Add path monitoring dashboard 1548
Configure policy-based routing policy 1548
Configure policy-based routing 1551
Configure PBR with path monitoring 1556
Useful CLIs for monitoring PBR 1558
Troubleshoot PBR 1560

PART XIV **Network Policies** 1565

CHAPTER 58 **Network Address Translation** 1567
NAT functions 1567
NAT basics 1568

NAT terminology 1568
NAT types 1568
NAT in routed and transparent mode 1569
NAT in routed mode 1569
NAT in transparent mode or within a bridge group 1570
Auto NAT and Manual NAT 1571
Auto NAT 1571
Manual NAT 1571
Comparing Auto NAT and Manual NAT 1572
NAT rule order 1572
NAT interfaces 1574
NAT exemption 1575
Configuring routing for NAT 1575
Addresses on the same network as the mapped interface 1576
Addresses on a unique network 1576
The same address as the real address (identity NAT) 1576
Requirements and prerequisites for NAT policies 1577
Guidelines for NAT 1577
Firewall mode guidelines for NAT 1577
IPv6 NAT guidelines 1578
IPv6 NAT best practices 1578
NAT support for inspected protocols 1578
FQDN destination guidelines 1580
Additional guidelines for NAT 1581
Manage NAT policies 1583
Create NAT policies 1584
Configure NAT policy targets 1584
Configure NAT for threat defense 1585
Customize NAT rules for multiple devices 1587
Search and filter the NAT rule table 1589
Enable, disable, or delete multiple rules 1590
NAT templates 1590
Create a NAT rule for internal hosts 1591
Create a NAT rule for external hosts 1592

Dynamic NAT 1593
About dynamic NAT 1593
Dynamic NAT disadvantages and advantages 1594
Configure dynamic auto NAT 1594
Configure dynamic manual NAT 1596
Dynamic PAT 1598
About dynamic PAT 1598
Dynamic PAT disadvantages and advantages 1599
PAT pool object guidelines 1599
Configure dynamic auto PAT 1600
Configure dynamic manual PAT 1602
Configure PAT with port block allocation 1605
Static NAT 1608
About static NAT 1608
Configure static auto NAT 1612
Configure static manual NAT 1614
Identity NAT 1617
Configure identity auto NAT 1618
Configure identity manual NAT 1619
NAT rule properties for Firewall Threat Defense 1621
Interface objects NAT properties 1622
Translation properties for auto NAT 1623
Translation properties for manual NAT 1624
PAT pool NAT properties 1625
Advanced NAT properties 1626
Translating IPv6 networks 1627
NAT64/46: translating IPv6 addresses to IPv4 1628
NAT64/46 example: inside IPv6 network with outside IPv4 internet 1628
NAT64/46 example: inside IPv6 network with outside IPv4 internet and DNS translation 1630
NAT66: translating IPv6 addresses to different IPv6 addresses 1634
NAT66 example, static translation between networks 1634
NAT66 example, simple IPv6 interface PAT 1637
Monitoring NAT 1640
Examples for NAT 1641

Providing access to an inside web server (static auto NAT) 1641
Dynamic auto NAT for inside hosts and static NAT for an outside web server 1643
Configure inside load balancer with multiple mapped addresses 1650
Single address for FTP, HTTP, and SMTP (static auto NAT with port translation) 1652
Different translation depending on the destination (dynamic manual PAT) 1660
Configure dynamic manual PAT with different destination ports 1666
NAT and site-to-site VPN 1672
DNS query and response rewriting using NAT 1677
DNS64 reply modification 1678
Create DNS reply modification with DNS server on outside 1685
Configure DNS reply modification with DNS server on host network 1688

CHAPTER 59 **Quality of Service** 1693
QoS 1693
QoS policies 1694
Requirements and prerequisites for QoS 1694
Guidelines and limitations for QoS policies 1695
Rate limiting with QoS policies 1695
Create a QoS policy 1696
Set target devices for a QoS policy 1696
Filter QoS rules by device 1697
Configuring QoS rules 1698
QoS rule components 1699
QoS rule conditions 1700
Interface rule conditions 1700
Network rule conditions 1701
User rule conditions 1701
Application rule conditions 1701
Best practice: use application-based rules over port-based rules 1703
Recommendation: use URL conditions for website access control 1704
Custom SGT rule conditions 1704
ISE SGT versus custom SGT rule conditions 1705
Autotransition from custom Security Group Tags to ISE Security Group Tags 1705

PART XV Secure Connections 1707

CHAPTER 60 Secure Connections Overview 1709

VPN types 1709
VPN basics 1710
Internet Key Exchange (IKE) 1711
IPsec 1712
VPN packet flow 1713
IPsec flow offload 1713
VPN licensing 1714
VPN encryption and performance 1714
Comply with security certification requirements 1714
Decide encryption algorithms for VPN policies 1715
Decide which hash algorithms to use 1715
Decide which Diffie-Hellman modulus group to use 1716
Decide VPN authentication methods 1717
Pre-shared keys 1717
PKI infrastructure and digital certificates 1718
Deprecated hash algorithms, encryption algorithms, and Diffie-Hellman modulus groups 1720
VPN topologies 1720
Point-to-point VPN topologies 1720
Hub and spoke VPN topology 1721
Full mesh VPN topology 1722
Implicit topologies 1723

CHAPTER 61 Site-to-Site VPN 1725

Site-to-site VPN 1725
Guidelines for configuring site-to-site VPN in Secure Firewall Threat Defense devices 1727
Limitations for configuring site-to-site VPN in Threat Defense devices 1727
Types of site-to-site VPN topologies 1728
License requirements for site-to-site VPN 1728
Prerequisites for site-to-site VPN 1729
Manage site-to-site VPNs 1730

Configure a policy-based site-to-site VPN 1731
Firewall Threat Defense VPN endpoint options 1732
Firewall Threat Defense VPN IKE options 1736
Firewall Threat Defense VPN IPsec options 1739
Firewall Threat Defense advanced site-to-site VPN deployment options 1741
Firewall Threat Defense VPN advanced IKE options 1741
Firewall Threat Defense VPN advanced IPsec options 1743
Firewall Threat Defense advanced site-to-site VPN tunnel options 1743
Configure virtual tunnel interfaces 1745
Virtual tunnel interfaces 1745
Static VTI 1746
Dynamic VTI 1747
Guidelines for virtual tunnel interfaces 1749
Limitations for Virtual Tunnel Interfaces 1752
Add a VTI interface 1752
Create a route-based site-to-site VPN 1754
Configure endpoints for a point-to-point VPN topology 1755
Configure endpoints for a hub and spoke topology 1757
Configure multiple hubs in a route-based VPN 1761
Route traffic through a backup VTI tunnel 1764
Configure dynamic VTI for a route-based site-to-site VPN 1765
Configure a virtual router with dynamic VTI 1765
Configure routing and AC policies for VTI 1766
View virtual tunnel interface details 1769
Deploy a SASE tunnel on Umbrella 1769
Guidelines for configuring SASE tunnels on Umbrella 1769
Limitations for configuring SASE tunnels on Umbrella 1770
Deploy a SASE tunnel on Umbrella 1770
Prerequisites for configuring Umbrella SASE tunnels 1771
Map Management Center Umbrella parameters and Cisco Umbrella API keys 1771
Configure a SASE tunnel for Umbrella 1773
Secure traffic using Cisco Secure Access and Firewall Threat Defense devices 1776
Licenses for automatic tunnel between Secure Access and Firewall Threat Defense devices 1776
Prerequisites for an automatic tunnel between Secure Access and Firewall Threat Defense devices 1777

Guidelines for an automatic tunnel between Secure Access and Firewall Threat Defense devices 1777
Limitation for automatic tunnel between Secure Access and Firewall Threat Defense devices 1777
Configure automatic tunnel between Secure Access and Firewall Threat Defense devices using SASE wizard 1778
Validate Secure Access integration with Firewall Threat Defense devices 1780
Monitor site-to-site topologies 1781
Monitor site-to-site VPNs using Site-to-Site VPN summary page 1781
Monitor site-to-site VPNs using Site-to-Site VPN dashboard 1781
Tunnel status widget 1782
Site-to-site VPN connection event monitoring 1785
View site-to-site VPN connection events 1786

CHAPTER 62 SD-WAN 1787
SD-WAN capabilities 1787
SD-WAN features 1788
SD-WAN wizard for secure branch network deployment 1789
Guidelines for using SD-WAN wizard 1790
Limitations for using SD-WAN wizard 1791
License requirements to use the SD-WAN wizard 1791
Prerequisites for using SD-WAN wizard 1792
Configure an SD-WAN topology using the SD-WAN wizard 1792
Add a dynamic virtual tunnel interface for a hub 1797
Sample configurations for dual-ISP deployment using SD-WAN wizard 1797
Dual ISP deployment: Two hubs and four spokes in the same region 1798
Dual ISP deployment: Two hubs and four spokes in different regions 1799
Verify tunnel statuses of an SD-WAN topology 1802
Monitor SD-WAN topologies 1805
SD-WAN summary dashboard 1805
Prerequisites for using SD-WAN summary dashboard 1806
Monitor WAN devices and interfaces using the SD-WAN summary dashboard 1807
Monitor application performance metrics of WAN interfaces using SD-WAN summary dashboard 1809

CHAPTER 63 Remote Access VPN 1811

Remote Access VPN Overview 1811
Remote Access VPN Features 1812
Secure Client Components 1814
Remote Access VPN Authentication 1814
Understanding Policy Enforcement of Permissions and Attributes 1816
Understanding AAA Server Connectivity 1816
License Requirements for Remote Access VPN 1817
Requirements and Prerequisites for Remote Access VPN 1817
Remote Access VPN Guidelines and Limitations 1818
Configuring a New Remote Access VPN Connection 1821
Prerequisites for Configuring Remote Access VPN 1822
Create a New Remote Access VPN Policy 1822
Update the Access Control Policy on the Secure Firewall Threat Defense Device 1825
(Optional) Configure NAT Exemption 1826
Configure DNS 1827
Add Secure Client Profile XML File 1827
(Optional) Configure Split Tunneling 1828
(Optional) Configure Dynamic Split Tunneling 1829
Verify Dynamic Split Tunneling Configuration 1830
Verify the Configuration 1830
Create a Copy of an Existing Remote Access VPN Policy 1831
Set Target Devices for a Remote Access VPN Policy 1831
Associate Local Realm with Remote Access VPN Policy 1832
Additional Remote Access VPN Configurations 1832
Configure Connection Profile Settings 1832
Configure IP Addresses for VPN Clients 1833
Configure AAA Settings for Remote Access VPN 1835
Create or Update Aliases for a Connection Profile 1850
Configure Access Interfaces for Remote Access VPN 1851
Configure Advanced Options for Remote Access VPN 1853
Cisco Secure Client Image 1853
Remote Access VPN Address Assignment Policy 1855
Configure Certificate to Connection Profile Mapping 1856
Configure Group Policies 1857

Configuring LDAP Attribute Mapping 1858
Configuring VPN Load Balancing 1859
Configure IPsec Settings 1862
Customize Cisco Secure Client 1868
Configure Secure Client Management VPN Tunnel 1881
Requirements and Prerequisites for Secure Client Management VPN Tunnel 1881
Limitations of Secure Client Management VPN Tunnel 1882
Configuring Secure Client Management VPN Tunnel on Firewall Threat Defense 1882
Multiple Certificate Authentication 1884
Guidelines and Limitations of Multiple Certificate Authentication 1884
Configuring Multiple Certificate Authentication 1885
Manage VPN Access of Remote Users Based on Geolocation 1886
Workflow to Manage VPN Access of Remote Users Based on Geolocation 1886
Guidelines and Limitations for Managing Remote Access VPN Users Based on Geolocation 1886
Monitor and Troubleshoot Service Access Policies 1887
Customizing Remote Access VPN AAA Settings 1890
Authenticate VPN Users via Client Certificates 1890
Configure VPN User Authentication via Client Certificate and AAA Server 1891
Manage Password Changes over VPN Sessions 1893
Send Accounting Records to the RADIUS Server 1894
Delegating Group Policy Selection to Authorization Server 1894
Override the Selection of Group Policy or Other Attributes by the Authorization Server 1895
Deny VPN Access to a User Group 1896
Restrict Connection Profile Selection for a User Group 1897
Update the Secure Client Profile for Remote Access VPN Clients 1898
RADIUS Dynamic Authorization 1899
Configuring RADIUS Dynamic Authorization 1899
Two-Factor Authentication 1900
Configuring RSA Two-Factor Authentication 1900
Configuring Duo Two-Factor Authentication 1901
Secondary Authentication 1903
Configure Remote Access VPN Secondary Authentication 1904
Single Sign-On Authentication with SAML 2.0 1906
Guidelines and Limitations for SAML 2.0 1906

Configuring a SAML Single Sign-On Authentication 1908
Configuring SAML Authorization 1909
Advanced Secure Client Configurations 1910
Configure Secure Client Modules on a Firewall Threat Defense 1910
Types of Secure Client Modules 1911
Prerequisites for Configuring Secure Client Modules 1912
Guidelines for Configuring Secure Client Modules 1913
Install Secure Client Modules using a Firewall Threat Defense 1914
Configure a Remote Access VPN Group Policy with Secure Client Modules 1914
Verify Secure Client Modules Configuration 1915
Configure Application-Based (Per App VPN) Remote Access VPN on Mobile Devices 1916
Prerequisites and Licensing for Configuring Per App VPN Tunnels 1916
Determine the Application IDs for Mobile Applications 1917
Configure Application-Based VPN Tunnels 1917
Verify Per App Configuration 1919
Remote Access VPN Examples 1919
How to Limit Secure Client Bandwidth Per User 1919
How to Use VPN Identity for User-Id Based Access Control Rules 1920
Configure Firewall Threat Defense Multiple Certificate Authentication 1920
Monitor Remote Access VPNs 1924
Remote Access VPN Dashboard 1924
VPN Session and User Information 1926
Viewing Remote Access VPN Active Sessions 1926
Viewing Remote Access VPN User Activity 1926

CHAPTER 64 **Dynamic Access Policies** 1927
Firewall Threat Defense dynamic access policies 1927
Hierarchy of policy enforcement of permissions and attributes in a Firewall Threat Defense device 1928
Prerequisites for dynamic access policy 1929
Guidelines and limitations for dynamic access policies 1929
Configure a Dynamic Access Policy (DAP) 1930
Create a dynamic access policy 1930
Create a dynamic access policy record 1930
Configure posture assessment criteria 1931

Configure AAA criteria settings for a DAP record 1932
Configure endpoint attribute selection criteria in DAP 1933
Add an anti-malware endpoint attribute to a DAP record 1933
Add a device endpoint attribute to a DAP record 1934
Add Secure Client endpoint attributes to a DAP record 1935
Add NAC endpoint attributes to a DAP record 1936
Add an application attribute to a DAP record 1936
Add a personal firewall endpoint attribute to a DAP record 1937
Add an operating system endpoint attribute to a DAP record 1937
Add a process endpoint attribute to a DAP 1938
Add a registry endpoint attribute to a DAP record 1939
Add a file endpoint attribute to a DAP record 1939
Add certificate authentication attributes to a DAP record 1940
Configure advanced settings for DAP 1941
Associate dynamic access policy with remote access VPN 1941

CHAPTER 65 **VPN Monitoring and Troubleshooting in Security Cloud Control** 1943
Monitor site-to-site VPNs using Site-to-Site VPN summary page 1943
Monitor Remote Access VPN Sessions 1943
SD-WAN summary dashboard 1943
Prerequisites for using SD-WAN summary dashboard 1944
Monitor WAN devices and interfaces using the SD-WAN summary dashboard 1945
Monitor application performance metrics of WAN interfaces using SD-WAN summary dashboard 1947

PART XVI **Zero Trust Network Access** 1949

CHAPTER 66 **Zero Trust Network Access Overview** 1951
Clientless and universal zero trust network access 1951
Zero trust network access 1952
Comparison of clientless and universal ZTNA 1952
History for zero trust network access 1953

CHAPTER 67 **Clientless Zero Trust Network Access** 1955
Zero Trust Access 1955

Components of a Clientless Zero Trust Network Access configuration 1956
How Clientless Zero Trust Network Access works 1957
How threat defense works with zero trust access 1958
Prerequisites for zero trust application policy 1960
Limitations for Clientless Zero Trust Network Access 1960
Create a zero trust application policy 1961
Create an application group 1963
Add certificate enrollment 1964
Create an application 1965
Create an ungrouped application 1967
Set targeted devices for zero trust access policy 1968
Edit a zero trust application policy 1969
Manage zero trust application policies 1970
Monitor zero trust sessions 1971

CHAPTER 68 **Universal Zero Trust Network Access** 1975
Universal Zero Trust Network Access 1975
Threat Defense with Universal ZTNA 1976
Prerequisites for Universal Zero Trust Network Access 1977
Restrictions for Universal Zero Trust Network Access 1978
Enable Cloud-Delivered Firewall Management Center in security cloud control 1979
Configure Security Devices 1979
Configure network connections 1985
Related documentation 1993

PART XVII **Access Control Policy Basics** 1995

CHAPTER 69 **Access Control Policies** 1997
Access control policies 1997
Access control policy components 1998
Access control policy default actions 1998
Access control policy inheritance 2000
Rule and other policy warnings 2001
Requirements and prerequisites for access control policies 2002

Manage access control policies 2002
Creating a basic access control policy 2004
Editing an access control policy 2004
Lock an access control policy 2006
Merge concurrent edits to an access control policy 2008
Configure the access control default action 2010
Manage access control policy inheritance 2011
Choose a base access control policy 2011
Lock settings in descendant access control policies 2012
Inherit access control policy settings from the base policy 2012
Enforce an access control policy in a domain 2013
Assign access control policy to devices 2014
Logging settings for access control policies 2014
Associate prefilter, decryption, and identity policies with an access control 2016
Identify and fix anomalies with policy analyzer & optimizer 2017
View rule hit counts 2018
Analyzing rule conflicts and warnings 2020
Search for rules 2022
History for access control policies 2023

CHAPTER 70 **Security Intelligence** 2027
Security intelligence 2027
License requirements for security intelligence 2028
Requirements and prerequisites for security intelligence 2028
Best practices for security intelligence 2028
Guidelines and limitations for security intelligence 2029
Security intelligence sources 2029
Configure security intelligence 2030
Security intelligence options 2032
Security intelligence categories 2034
Block list icons 2036
Configuration example: Security intelligence blocking 2036
Security intelligence monitoring 2037
Override security intelligence blocking 2038

Troubleshooting security intelligence 2039
Security intelligence categories missing from the Available Options List 2039

CHAPTER 71 Access Control Rules 2041
Access control rules 2041
Access control rule management 2043
Access control rule components 2043
Access control rule order 2045
Access control rule actions 2046
Access control rule monitor action 2046
Access control rule trust action 2046
Access control rule blocking actions 2047
Access control rule interactive blocking actions 2047
Access control rule allow action 2048
Deep inspection using file and intrusion policies 2049
File and intrusion inspection order 2050
How access control traffic handling works with intrusion and file policies 2051
Requirements and prerequisites for access control rules 2052
Guidelines and limitations for access control rules 2053
Best practices for application control 2054
Recommendations for application control 2054
Best practice for choosing between application matching and port matching 2056
Application-specific requirements for filtering 2057
Best practices for access control rules 2058
General best practice for access control 2058
Best practices for ordering rules 2059
Application rule order 2060
Rule preemption 2060
Best practice for rule actions and rule order 2061
Best practices for simplifying and focusing rules 2061
Maximum number of access control rules and intrusion policies 2062
Access control rule management 2062
Add an access control rule category 2062
Introduction to real-time Policy Analyzer and Optimizer 2063

How real-time Policy Analyzer and Optimizer works 2064
Prerequisites for using real-time Policy Analyzer and Optimizer 2064
Limitations of real-time Policy Analyzer and Optimizer 2064
Manage rule anomalies 2065
Create and edit access control rules 2066
Access control rule conditions 2068
Enable and disable access control rules 2076
Copy access control rules from one access control policy to another 2077
Move access control rules to a prefilter policy 2077
Position an access control rule 2079
Add comments to an access control rule 2080
Examples for access control rules 2080
How to control access using security zones 2080
Control application usage 2081
Block threats 2082

CHAPTER 72 **URL Filtering Rules** 2085
URL filtering 2085
URL filtering with category and reputation 2086
URL category and reputation descriptions 2087
URL filtering data from the Cisco cloud 2087
Best practices for URL filtering 2088
HTTPS traffic filtering 2091
Use categories in URL filtering 2092
QUIC Fingerprinting 2093
License requirements for URL filtering 2093
Requirements and prerequisites for URL filtering 2094
Configure URL filtering with category and reputation 2094
Enable URL filtering using category and reputation 2096
URL filtering options 2096
Configure URL conditions 2097
Rules with URL conditions 2099
URL rule order 2099
DNS filtering 2099

Enable DNS filtering to identify URLs during domain lookup 2100
DNS filtering behavior for block actions 2101
DNS filtering and events 2101
Manual URL filtering 2101
Manual URL filtering options 2102
Supplement or selectively override category and reputation-based URL filtering 2102
HTTP response pages 2103
Limitations to HTTP response pages 2103
Requirements and prerequisites for HTTP response pages 2104
Choose HTTP response pages 2105
Interactive blocking with HTTP response pages 2105
Configure interactive blocking 2106
Set the user bypass timeout for a blocked website 2107
Configure URL filtering health monitors 2107
Dispute URL category and reputation 2107
Take action if the URL category set changes 2108
URL category and reputation changes 2109
Troubleshoot URL filtering 2110

PART XVIII **Decryption Policies and Encrypted Visibility for Access Control** 2113

CHAPTER 73 **Traffic Decryption Overview** 2115
Traffic decryption 2115
How TLS/SSL handshakes work 2116
ClientHello message handling 2117
How ServerHello and server certificate messages work 2120
Decryption rule and policy basics 2121
Guidelines for traffic decryption 2122
Best practice for deciding when to decrypt traffic 2122
Decryption policy types 2123
Decrypt and re-sign (outgoing traffic) 2125
Incoming traffic decryption 2125
Known key decryption for incoming traffic 2126
Other decryption rule actions 2126

Decryption rule components 2126
Rule-based decryption rule order evaluation 2127
Multi-rule example 2128
TLS crypto acceleration 2130
TLS crypto acceleration guidelines and limitations 2131
View the status of TLS crypto acceleration 2132
History for decryption policy 2133

CHAPTER 74 Standard Decryption Policies 2137
Standard decryption policies 2137
Decryption policy types 2137
Standard decryption policy deployment issues to versions earlier than 10.0.0 2139
Create a new standard decryption policy 2139
Create a standard decryption policy with inbound protection 2139
Security zones 2140
Internal server details (inbound decryption) 2142
Create a standard decryption policy with outbound protection 2144
Security zones 2145
Decrypt networks and users 2146
Internal certificate authority 2148
Bypass traffic when decrypting 2150
Block connections 2155
Standard decryption policy advanced options 2157
Add trusted CA certificates 2159
Decryption policy actions 2159
Convert a standard decryption policy to a rule-based decryption policy 2160
Copy a decryption policy 2160
Generate a decryption policy report 2161

CHAPTER 75 Rule-Based Decryption Policies 2163
Rule-based decryption policies 2163
Decryption policy types 2163
Create a rule-based decryption policy 2165
Create a rule-based decryption policy with outbound connection protection 2165

Create a rule-based decryption policy with inbound connection protection 2169
Rule-based decryption policy block connections 2171
Configure Rule-based decryption policy exclusions 2171
Generate an internal CA for outbound protection 2175
Upload an internal CA for outbound protection 2176
Upload an internal certificate for inbound protection 2177
Install an internal CA on client machines 2177
Create a rule-based decryption policy with other rule actions 2178
Rule-based decryption policy default actions 2179
Default handling options for undecryptable traffic 2180
Set default handling for undecryptable traffic 2181
Rule-based decryption policy advanced options 2182
TLS 1.3 decryption best practices 2184

CHAPTER 76 **Rule-Based Decryption Rules** 2187
Rule-based decryption 2187
Requirements and prerequisites for Rule-based decryption rules 2187
Rule-based decryption rules guidelines and limitations 2188
Guidelines for using TLS/SSL decryption 2188
Rule-based decryption rules unsupported features 2190
Guidelines for TLS/SSL Do Not Decrypt configuration 2190
Best practices for TLS/SSL incoming traffic decryption 2191
TLS/SSL Decrypt - Resign guidelines 2192
TLS/SSL block guidelines 2194
TLS/SSL certificate pinning guidelines 2194
TLS/SSL heartbeat guidelines 2195
TLS/SSL anonymous cipher suite limitation 2195
TLS/SSL normalizer guidelines 2195
Other Rule-based decryption rule guidelines 2195
Best practice for Decryption policy rule order 2196
Rule-based decryption rule traffic handling 2197
Encrypted traffic inspection configuration 2199
Rule-based decryption rule order evaluation 2200
Rule-based decryption rule conditions 2201

Security zone rule conditions 2202
Network rule conditions 2203
VLAN tag rule conditions 2203
User rule conditions 2204
Application rule conditions 2204
Best practice: use application-based rules over port-based rules 2206
Category rule conditions 2206
Server certificate-based rule conditions 2207
Certificate rule conditions 2207
Conditions for distinguished name rules in decryption policies 2208
Trust external certificate authorities 2213
Certificate Status Decryption rule conditions 2214
Cipher suite rule conditions 2217
Encryption protocol version rule conditions 2220
Rule-based decryption rule actions 2220
Decryption rule monitor action 2220
Decryption rule Do Not Decrypt action 2221
Decryption rule blocking actions 2222
Decryption rule decrypt actions 2222
Incoming traffic decryption actions 2223
Decrypt - Resign rule action 2226
Select internal certificate objects 2226
Monitor TLS/SSL hardware acceleration 2226
Informational counters 2226
Alert counters 2227
Error counters 2227
Fatal counters 2227
Troubleshoot Rule-based decryption rules 2228
TLS/SSL oversubscription 2228
Troubleshoot TLS/SSL oversubscription 2229
TLS heartbeat 2229
Troubleshoot TLS heartbeat 2230
TLS/SSL pinning 2231
Troubleshoot TLS/SSL pinning 2231

Troubleshoot unknown or bad certificates or certificate authorities 2233
Verify TLS/SSL cipher suites 2235
Troubleshoot using crypto archives 2236

CHAPTER 77 **Encrypted Visibility Engine** 2239
Encrypted Visibility Engine 2239
How EVE works 2240
Indications of compromise events 2241
QUIC fingerprinting in EVE 2241
Configure EVE 2242
View Encrypted Visibility Engine events 2242
Configure EVE exception rules 2243
Add exception rule from unified events 2245
Upgrade EVE exception rules 2245
Use case - Block Traffic Based on the EVE Threat Confidence Score 2247
About Encrypted Visibility Engine 2247
Benefits 2247
Sample Business Scenario 2247
Prerequisites 2247
High-Level Workflow 2247
Configure Block Thresholds in EVE 2248
View EVE Events 2251
Additional References 2252

PART XIX **Identity Policies for Access Control** 2253

CHAPTER 78 **Identity Overview** 2255
User identity 2255
Identity terminology 2255
About user identity sources 2256
Best practices for user identity 2258
Identity deployments 2260
Set up an identity policy 2263
User activity database 2267

Users database 2267
Recommendation: identity realm limit 2269
Cloud-Delivered Firewall Management Center host and user limits 2269
Cloud-Delivered Firewall Management Center host limit 2269
Cloud-Delivered Firewall Management Center user limit 2270
User limits for microsoft azure active directory realms 2271

CHAPTER 79 Realms 2273
Realms and realm sequences 2273
Realms and trusted domains 2275
Supported servers for realms 2278
Supported server object class and attribute names 2279
License requirements for realms 2280
Requirements and prerequisites for realms 2281
Prerequisites for kerberos authentication 2281
Realm fields 2281
Realm directory and synchronize fields 2285
Create a proxy sequence 2288
Create a Microsoft Azure AD (SAML) realm 2289
Create a Microsoft Azure AD realm for passive authentication 2290
How Entra ID and Cisco ISE authentication with resource owned password credentials works 2291
How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS 2292
Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD 2293
Configure Microsoft Entra ID for passive authentication 2293
Configure Entra ID basic settings 2295
Get required information For Your Microsoft Azure AD realm 2296
Create a Microsoft Azure AD (SAML) realm for passive authentication 2298
Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal) 2302
Configure Entra ID basic settings 2304
Configure a single sign-on (SSO) app in Entra ID 2305
Create a decryption rule with decrypt - resign action 2307
Get required information for your Microsoft Azure AD realm (active authentication only) 2307
Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal) 2310
Create a Microsoft Azure AD Realm 2316

How Entra ID and Cisco ISE authentication with resource owned password credentials works 2316
How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS 2317
Create a Microsoft Azure AD realm for passive authentication 2318
Configure Microsoft Entra ID for passive authentication 2319
Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD 2320
Get required information For Your Microsoft Azure AD realm 2321
Create an Azure AD realm 2323
Azure AD user session timeout 2325
Create an LDAP realm or an Active Directory realm and realm directory 2325
Configure the Realm 2327
Connect securely to Active Directory or LDAP 2328
Find the active directory server's name 2328
Export the active directory server's root certificate 2329
Synchronize users and groups 2330
Create a realm sequence 2332
Cross-domain trust setups with Cloud-Delivered Firewall Management Center 2333
Configure Security Cloud Control for cross-domain-trust step 1: Configure realms and directories 2334
Configure the Security Cloud Control for cross-domain trust Step 2: Synchronize users and groups 2339
Configure the Security Cloud Control for cross-domain trust step 3: Resolve issues 2340
Manage a realm 2341
Compare realms 2342
Troubleshoot realms and user downloads 2342
Troubleshoot cross-domain trust 2345
History for realms 2348

CHAPTER 80 **Identity Source: Captive Portal** 2351
Captive portal identity source 2351
Hostname redirect 2352
License requirements for captive portal 2352
Requirements and prerequisites for captive portal 2352
Best practices for captive portal guidelines and limitations 2353
Configure the captive portal for user control 2355
Create a network object 2357
Create an identity policy and active authentication rule 2358

Update a custom authentication form 2360
Create a TCP port access control rule 2360
Create a user access control rule 2362
Create a decryption policy with an outbound rule 2362
Create A decryption policy with a Decrypt-Resign rule 2365
Associate Identity and decryption policies with the access control policy 2366
Captive portal fields 2367
Exclude applications from captive portal 2368
Troubleshoot the captive portal identity source 2369
History for captive portal 2370

CHAPTER 81 **Identity Source: ISE/ISE-PIC** 2373
The ISE/ISE-PIC identity source 2373
Source and destination security group tag (SGT) matching 2374
License requirements for ISE/ISE-PIC 2375
Requirements and prerequisites for ISE/ISE-PIC 2375
Best practices for ISE/ISE-PIC integration 2376
ISE/ISE-PIC configurations for user control 2379
Configure ISE/ISE-PIC without a realm 2379
Configure ISE/ISE-PIC for user control using a realm 2380
ISE/ISE-PIC configuration 2382
Configure security groups and SXP publishing in ISE 2382
Export certificates from the ISE/ISE-PIC server for use in the Cloud-Delivered Firewall Management Center 2384
Export a system certificate 2385
Generate a self-signed certificate 2386
Import ISE/ISE-PIC certificates 2387
Configure ISE for user control 2387
ISE/ISE-PIC configuration fields 2389
Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source 2390
Cisco ISE quick configuration 2391
Configure prerequisites for Cisco ISE quick configuration 2391
Quick configuration 2392
Cisco Identity Services Engine (Cisco ISE) quick configuration results 2394

Configure Cisco ISE advanced configuration 2396
ISE/ISE-PIC configuration fields 2398
Troubleshoot Cisco ISE/ISE-PIC or Cisco TrustSec issues 2399
History for ISE/ISE-PIC 2401

CHAPTER 82 **Identity Source: Passive Identity Agent** 2403
The passive identity agent identity source 2403
Deploy the passive identity agent 2404
Simple passive identity agent deployment 2405
Single passive identity agent monitoring multiple domain controllers 2405
Multiple passive identity agents monitoring multiple domain controllers 2406
Configure Passive identity agent primary and secondary agent deployments 2408
How to configure a passive identity agent identity source 2409
Configure the passive identity agent 2410
Enable the dynamic attributes connector 2411
Create a Microsoft Active Directory realm 2411
Create a passive identity agent identity source 2411
Create a standalone passive identity agent identity source 2412
Create a primary or secondary passive identity agent identity source 2414
Passive identity agent roles 2417
Get an API token for the passive identity agent 2418
Passive identity agent installation 2418
Prerequisites to install the passive identity agent 2419
Install the passive identity agent software 2422
Uninstall the passive identity agent software 2425
Upgrade the passive identity agent software 2426
Monitor the passive identity agent 2426
Manage the passive identity agent 2427
Edit passive identity agents 2428
Delete a standalone passive identity agent 2428
Delete primary and secondary passive identity agents 2428
Troubleshoot the passive identity agent 2429
Security requirements for the passive identity agent 2430
Internet access requirements for the passive identity agent 2431

History for the passive identity agent 2431

CHAPTER 83 **Identity Source: pxGrid Cloud Identity (ISE 3.3 and Earlier)** 2433

pxGrid Cloud identity source 2433

Restrictions for the pxGrid Cloud identity source 2434

How the pxGrid Cloud identity source works 2434

pxGrid Cloud identity source configuration 2436

Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier) 2436

Enable the pxGrid Cloud service in Cisco ISE 2439

Register Cisco ISE with the Catalyst Cloud Portal 2439

Register the pxGrid Cloud connection with Cisco ISE 2442

Create a pxGrid Cloud identity source 2443

Create an app instance 2443

Create the identity source 2445

Activate the app instance 2446

Activate the pxGrid Cloud identity source 2448

Test the pxGrid Cloud identity source 2450

Create dynamic attributes filters 2455

Create access control rules or DNS rules using dynamic attributes filters 2456

Troubleshoot the pxGrid Cloud identity source 2457

Primary device cannot be processed 2458

Deactivating and deleting the pxGrid Cloud identity source 2459

Deactivate the pxGrid Cloud app instance 2459

Delete the pxGrid Cloud identity source 2462

CHAPTER 84 **Identity Source: pxGrid Cloud Identity (ISE 3.4 and Later)** 2465

pxGrid Cloud identity source 2465

Restrictions for the pxGrid Cloud identity source 2466

How the pxGrid Cloud identity source works 2466

pxGrid Cloud identity source configuration 2468

Configuring a pxGrid Cloud identity source 2468

Enable the pxGrid Cloud service in Cisco ISE 2470

Verify Cisco ISE registration in the Catalyst Cloud Portal 2472

Create an app instance 2473

Create the identity source 2474
Activate the app instance 2475
Activate the pxGrid Cloud identity source 2478
Test the pxGrid Cloud identity source 2480
Troubleshoot the pxGrid Cloud identity source 2485
Primary device cannot be processed 2485
Create dynamic attributes filters 2486
Create access control rules or DNS rules using dynamic attributes filters 2487
Deactivating and deleting the pxGrid Cloud identity source 2489
Deactivate the pxGrid Cloud app instance 2489
Delete the pxGrid Cloud identity source 2492

CHAPTER 85 **Identity Source: Remote Access VPN** 2495
Remote access VPN identity sources 2495
Configure Remote Access VPN for AD-authenticated users 2496
Hybrid Entra ID and on-premises AD deployments 2497
Troubleshooting guidelines for remote access VPN identity source 2497
Not Observing Correct Settings for VPN Statistics 2498

CHAPTER 86 **Identity source: TS agent** 2501
Terminal services agent identity source 2501
Best practices for TS agent configuration 2502
Identity source: TS agent 2502
Troubleshoot the TS agent identity source 2502
History for TS agent 2503

CHAPTER 87 **Identity Policies** 2505
Identity policies 2505
License requirements for identity policies 2506
Requirements and prerequisites for identity policies 2506
Create an identity policy 2507
Create an identity mapping filter 2508
Identity rule conditions 2509
Security zone rule conditions 2509

Network rule conditions 2510
Redirect to host name network rule conditions 2510
VLAN tag rule conditions 2511
Best practice: use application-based rules over port-based rules 2512
Port, protocol, and ICMP code rule conditions (concept) 2512
Realm & settings rule conditions 2513
Create an identity rule 2515
Identity rule fields 2517
Sample identity policies and rules 2518
Create an identity policy with a passive authentication rule 2519
Create a sample identity policy with an active authentication rule 2520
Active authentication using a realm 2522
Active authentication using a realm sequence 2523
Manage an identity policy 2524
Manage an identity rule 2525
Troubleshoot user control 2526

PART XX **Advanced Policies and Settings for Access Control** 2529

CHAPTER 88 **Prefilter Policies** 2531
Prefilter policies 2531
Default prefilter policy 2532
Prefilter policy rule order 2532
Tunnel vs prefilter rules 2533
How the system processes plain-text tunnels 2533
Using tunnel zones to apply access control at the tunnel level 2534
Large flow offloads 2535
Restrictions for flow offload implementation 2537
Prefilter vs access control policy 2539
Requirements and prerequisites for prefilter policies 2540
Guidelines and limitations for the prefilter policy 2541
Guidelines and limitations for encapsulated traffic handling 2541
Best practices for fastpath prefiltering 2542
Configure prefilter policies 2543

Configuring the default action 2545
Configure prefilter rules 2545
VLAN tags rule conditions 2547
Port rule conditions for prefilter rules 2547
Configure tunnel rules 2548
Creating tunnel zones 2550
Move prefilter rules to an access control policy 2551
Examples of prefilter policy 2552
Rezone tunnels for customized inspection 2552
Offload large flows 2554

CHAPTER 89 **DNS Policies for Security Intelligence** 2557
DNS policies 2557
Cisco umbrella DNS policies 2558
DNS policy components 2559
License requirements for DNS policies 2560
Requirements and prerequisites for DNS policies 2560
Managing DNS and Umbrella DNS Policies 2560
Create basic DNS policies 2561
Edit DNS policies 2562
DNS rules 2562
Create and edit DNS rules 2563
DNS rule management 2564
Enable and disable DNS rules 2564
DNS rule order evaluation 2565
DNS rule actions 2565
DNS rule conditions 2566
Security zone rule conditions 2567
Network rule conditions 2567
VLAN tag rule conditions 2568
DNS policy rule conditions 2568
Dynamic attributes rule conditions 2569
DNS rules 2569
Control traffic based on DNS and security zone 2569

Control traffic based on DNS and network 2570
Control traffic based on DNS and VLAN 2571
Control traffic based on DNS list or feed 2572
Controlling traffic based on security group tag or dynamic attributes 2572
Use a DNS sinkhole to enforce content restriction 2573
Requirement: deploy DNS policy configuration changes 2575
Cisco umbrella DNS policies 2575
Redirect DNS requests to Cisco umbrella 2576
Prerequisites for configure the umbrella DNS connector 2576
Configure Cisco umbrella connection settings 2578
Create an umbrella DNS policy 2579
Edit Umbrella DNS policies and rules 2579
Associate the umbrella DNS policy with an access control policy 2580

CHAPTER 90 File Policies for Network Malware Protection 2583
Network malware protection and file policies 2583
File policies 2583
Prerequisites for file policies 2584
License requirements for file and malware policies 2585
Best practices for file policies and malware detection 2585
File rule best practices 2585
File detection best practices 2586
File blocking best practices 2587
File policy best practices 2588
Configure malware protection 2589
Plan and prepare for malware protection 2590
Configure file policies 2591
Add file policies to your access control configuration 2591
Configure an access control rule to perform malware protection 2592
Verify and monitor malware protection 2593
Cloud connections for malware protection 2594
AMP cloud connection configurations 2594
Change AMP options 2594
Dynamic analysis connections 2595

Prerequisites for dynamic analysis 2595
View the default dynamic analysis connection 2595
Enable access to dynamic analysis results in the public cloud 2596
Maintain your system: file types eligible for dynamic analysis 2596
File policies and file rules 2597
Create or edit a file policy 2597
Advanced and archive file inspection options 2598
Manage file policies 2601
File rules 2602
File rule components 2603
File rule actions 2604
Create file rules 2611
Access control rule logging for malware protection 2613
Retrospective disposition changes 2613
File and malware inspection performance and storage options 2613
Tune file and malware inspection performance and storage 2615
(Optional) Malware protection with Secure Endpoint 2615
Comparison of malware protection: Firepower vs. Secure Endpoint 2616
Integrate Secure Firewall with Secure Endpoint 2617
Benefits of integrating Firepower and Secure Endpoint 2617
Secure Endpoint and AMP private cloud 2618
Integrate Firepower and Secure Endpoint 2618

CHAPTER 91 **Dynamic Attributes Connector** 2621
About the Dynamic Attributes Connector 2621
How It Works 2623
Enable the dynamic attributes connector 2624
Dashboard 2624
Dashboard of an unconfigured system 2625
Dashboard of a configured system 2626
Add, edit, or delete connectors 2627
Add, edit, or delete dynamic attributes filters 2629
Connectors 2631
Amazon Web Services connector user permissions and imported data 2631

Create an AWS user with minimal permissions for the dynamic attributes connector 2632
Create an AWS connector 2633
Amazon Web Services security groups connector—About user permissions 2634
Create an AWS security groups connector 2634
Create an AWS service tags connector 2635
Azure connector—About user permissions and imported data 2636
Create an Azure user with minimal permissions for the dynamic attributes connector 2637
Create an Azure connector 2639
Create an Azure service tags connector 2640
Create a Multicloud Defense connector 2641
Create a Cisco cyber vision connector 2643
Create a generic text connector 2645
Manually get a certificate authority (CA) chain 2646
Create a GitHub connector 2649
Google cloud connector user permissions and imported data 2650
Create a Google Cloud user with minimal permissions for the dynamic attributes connector 2651
Create a google cloud connector 2652
Create an office 365 connector 2653
Tenable connector 2654
About the Tenable connector 2654
Get the Tenable API key and secret 2655
Create a Tenable connector 2655
About Tenable dynamic objects in IDS, IPS, and access control policies 2657
Create a Webex connector 2657
Create a Zoom connector 2658
Create an adapter 2659
How to create an On-Premises Firewall Management Center adapter 2660
How to create a Cloud-Delivered Firewall Management Center adapter 2661
Create dynamic attributes filters 2661
Dynamic attribute filter examples 2663
Disable the dynamic attributes connector 2664
Dynamic firewall 2664
Dynamic Firewall 2664
How to configure the dynamic firewall 2665

Get required information for Identity Intelligence 2667
Create an identity source and realm for the dynamic firewall 2669
Create a dynamic firewall instance 2669
Associate an identity source with Identity Intelligence 2670
Configure Identity Intelligence 2671
View system-defined filters 2672
View system-defined access control rules 2673
Edit the user exclusion list 2675
View and edit the system-created access control policy 2676
Create dynamic attributes filters 2677
Use Dynamic Objects in Access Control Policies 2679
Dynamic objects in access control rules 2679
Create dynamic attributes filters 2679
Dynamic attributes rule conditions 2681
Create access control rules or DNS rules using dynamic attributes filters 2681
Use dynamic objects in DNS policies 2683
Troubleshoot the Dynamic Attributes Connector 2683
Troubleshoot error messages 2684
Get your tenant ID 2684

CHAPTER 92 **Advanced Settings for Access Control** 2687
Requirements and prerequisites for advanced settings 2687
Configure advanced settings for the access control policy 2688
General settings 2691
TLS server identity discovery 2692
Intelligent application bypass 2693
Configure intelligent application bypass 2694
IAB logging and analysis 2696
Transport/network layer preprocessor settings 2699
Detection enhancement settings 2700
Adaptive profile updates 2701
Adaptive profile updates and recommended rules 2702
Performance settings 2702
Latency-based performance settings 2704

CHAPTER 93 **Service Policies** 2707

Threat defense service policies 2707

Service policy relationships with FlexConfig and other features 2708

Connection settings 2708

Requirements and prerequisites for service policies 2709

Guidelines and limitations for service policies 2709

Configure threat defense service policies 2710

Configure a service policy rule 2711

Bypass TCP state checks for asymmetrical routing (TCP state bypass) 2714

The asymmetrical routing problem 2714

Guidelines and limitations for TCP state bypass 2715

Configure TCP state bypass 2716

Disable TCP sequence randomization 2718

Monitor service policies 2719

Examples for service policy rules 2720

Protect servers from a SYN flood DoS attack (TCP intercept) 2720

Configure the Firewall Threat Defense device to appear on traceroutes 2723

CHAPTER 94 **Threat Detection** 2725

Portscan detection and prevention 2725

Pre-defined sensitivity levels for portscan detection 2725

Detection in the low sensitivity level 2727

Requirements and prerequisites for threat detection 2727

Limitations for threat detection configuration 2728

Best practices for portscan prevention 2729

Configure portscan detection and prevention 2729

Monitoring threat detection 2731

View portscan alerts 2731

Monitor portscan on the firewall 2732

Unblocking a host 2733

CHAPTER 95 **Elephant Flow Detection** 2735

Elephant flow detection and remediation 2735

Elephant flow upgrade from Intelligent Application Bypass 2736
Configure elephant flow detection 2736
Use case - Configure Elephant Flow Detection Outcomes 2739
Elephant flows 2739
Benefits of elephant flow detection and remediation 2739
How elephant flows work 2739
Sample business scenarios 2740
Prerequisites 2741
Configure elephant flow parameters 2741
View events for elephant flows 2743
Configure elephant flow remediation exemption 2744
View events for elephant flow remediation exemption 2747
Additional references 2747

CHAPTER 96 **Policy Analyzer and Optimizer** 2749
About Policy Analyzer and Optimizer 2749
Common Policy Analyzer and Optimizer concepts 2750
Prerequisites to use Policy Analyzer and Optimizer 2751
Enable Policy Analyzer and Optimizer for Security Cloud Control-managed On-Premises Firewall Management Center 2752
Policy Analyzer and Optimizer licensing requirements 2753
Open Policy Analyzer and Optimizer and select a data source 2753
Sync and analysis behavior 2754
Search and filter policies in Policy Analyzer and Optimizer 2754
Access control policy analysis and optimization 2754
Analyze access control policies 2755
Access control policy analysis summary 2756
Analyze duplicate rules in access control policy 2757
Analyze expired rules in access control policy 2759
Analyze mergeable rules in access control policy 2759
Analyze overlapping objects in access control policy 2760
Analyze policy insights in access control policy 2761
Access control policy remediation 2762
Download access control analysis report 2764

Network Address Translation policy analysis 2764
Analyze NAT policies 2765
NAT rule sections and comparison behavior 2765
Supported NAT address families for Policy Analyzer and Optimizer 2766
NAT policy analysis summary 2766
Download NAT analysis report 2767
NAT policy remediation 2768
NAT remediation history 2769
NAT Policy Analyzer and Optimizer and AgenticOps insights 2770
Troubleshooting Policy Analyzer and Optimizer 2770
Policy Analyzer and Optimizer Does Not Analyze Policies 2771
Policy Analyzer and Optimizer Does Not Fetch Policies 2771
Frequently Asked Questions About Policy Analyzer and Optimizer 2771

PART XXI **Custom Intrusion Policies for Access Control** 2773

CHAPTER 97 **Intrusion Prevention** 2775
Intrusion prevention 2775
Snort inspection engine 2777
Snort 3 2777
Snort 3 multi-process support 2780
Disable snort 3 multi-process support 2780
Multi-layer inspection 2781
Header inspection 2781
Layer 7 features 2782
Example scenarios 2782
Encapsulation layer limits and inspection behavior 2782
Supported protocols and inspection capabilities of LINA and Snort 3 2783
Guidelines and limitations for network analysis and intrusion policies 2784
Prerequisites for network analysis and intrusion policies 2784
How policies examine traffic for intrusions 2785
Decoding, normalizing, and preprocessing: Network analysis policies 2786
Access control rules: Intrusion policy selection 2787
Intrusion inspection: Intrusion policies, rules, and variable sets 2788

Intrusion event generation 2790
System-provided and custom network analysis and intrusion policies 2790
System-provided network analysis and intrusion policies 2791
Benefits of custom network analysis and intrusion policies 2793
Custom network analysis policy benefits 2793
Benefits of custom intrusion policies 2794
Limitations of custom policies 2795

CHAPTER 98 **Migrate from Snort 2 to Snort 3** 2799
Snort 3 inspection engine 2799
Snort 2 versus Snort 3 2800
Migrating from Snort 2 to Snort 3 2800
Prerequisites for migrating from Snort 2 to Snort 3 2801
Enable Snort 3 on an individual device 2801
Enable Snort 3 on multiple devices 2802
Convert Snort 2 Custom IPS Rules to Snort 3 2802
Convert all Snort 2 custom rules across all intrusion policies to Snort 3 2803
Convert Snort 2 custom rules of a single intrusion policy to Snort 3 2804
View Snort 2 and Snort 3 base policy mapping 2804
Synchronize Snort 2 rules with Snort 3 2805
Deploy configuration changes 2806
Examples for migration 2808
Migrate from Snort 2 to Snort 3 2808
Benefits of migrating to snort 3 2808
Sample business scenario 2809
Best practices for migrating from Snort 2 to Snort 3 2809
Prerequisites for migration 2810
End-to-end migration workflow 2810
Enable Snort 3 on Threat Defense device 2811
Convert Snort 2 rules of a single intrusion policy to Snort 3 2812

CHAPTER 99 **Intrusion Policies** 2819
Intrusion policies 2819
Create a custom snort 3 intrusion policy 2820

Edit Snort 3 Intrusion Policies 2821
Rule group reporting 2825
Rule action logging 2825
Change the base policy of an intrusion policy 2826
View Snort 2 and Snort 3 base policy mapping 2827
Synchronize Snort 2 rules with Snort 3 2827
Manage intrusion policies 2829
Access control rule configuration to perform intrusion prevention 2830
Access control rule configuration and intrusion policies 2830
Configure an access control rule to perform intrusion prevention 2831
Tune Intrusion Policies using Rules 2831
Tuning intrusion rules 2831
Intrusion rule types 2832
Custom rules in Snort 3 2833
View Snort 3 intrusion rules in an intrusion policy 2836
Intrusion rule actions 2836
Intrusion rule action options 2837
Set intrusion rule action 2837
Intrusion event notification filters in an intrusion policy 2838
Intrusion event thresholds 2838
Intrusion policy suppression configuration 2841
Add intrusion rule comments 2843
Prepare custom rules for snort 3 conversion 2844
Convert all Snort 2 custom rules across all intrusion policies to Snort 3 2844
Convert Snort 2 custom rules of a single intrusion policy to Snort 3 2845
Add custom rules to rule groups 2846
Add rule groups with custom rules to an intrusion policy 2847
Manage custom rules in Snort 3 2847
Delete custom rules 2848
Delete rule groups 2849
Recommended rules 2850
Snort 3 rule changes in LSP updates 2850
Secure Firewall recommended rules 2850
Generate new Secure Firewall recommendations in Snort 3 2851

Mitigate threats using MITRE framework in Snort 3 intrusion policies 2854
MITRE framework 2854
MITRE framework benefits 2855
Sample business scenario for MITRE network 2855
Prerequisites for MITRE framework 2855
View and edit your Snort 3 intrusion policy 2856
View intrusion events 2860
Additional references 2863

CHAPTER 100 **Intrusion Prevention Performance Tuning** 2865
Intrusion prevention performance tuning 2865
License requirements for intrusion prevention performance tuning 2866
Requirements and prerequisites for intrusion prevention performance tuning 2866
Limiting pattern matching for intrusions 2866
Regular expression limits overrides for intrusion rules 2867
Override regular expression limits for intrusion rules 2868
Per packet intrusion event generation limits 2869
Limiting intrusion events generated per packet 2869
Packet and intrusion rule latency threshold configuration 2870
Latency-based performance settings 2871
Packet latency thresholding 2871
Packet latency thresholding notes 2873
Enable packet latency thresholding 2873
Configure packet latency thresholding 2873
Rule latency thresholding 2874
Rule latency thresholding notes 2876
Configure rule latency thresholding 2876
Intrusion performance statistic logging configuration 2877
Configure intrusion performance statistic logging 2878

CHAPTER 101 **Network Analysis Policies** 2879
Network analysis policies 2879
Manage network analysis policies 2880
Snort 3 definitions and terminologies for network analysis policies 2881

Custom network analysis policy creation for Snort 3 2883
Common Industrial Protocol safety 2887
Detect and block safety segments in CIP packets 2888
Network analysis policy mapping 2889
View network analysis policy mapping 2889
Create a network analysis policy 2889
Modify the network analysis policy 2890
Search for an inspector on the Network Analysis Policy window 2890
Copy the inspector configuration 2890
Customize the network analysis policy 2891
Make inline edit for an inspector to override configuration 2894
Revert unsaved changes during inline edits 2895
View the list of inspectors with overrides 2896
Revert overridden configuration to default configuration 2897
Validate Snort 3 policies 2897
Examples of custom network analysis policy configuration 2899
Network analysis policy settings and cached changes 2910
Generate Snort 3 recommendations 2911
Snort 3 rule recommendations 2911
Benefits 2911
Sample business scenario 2911
Best practices 2912
Prerequisites for generating recommendations 2912
Generate Snort 3 recommendations 2912

CHAPTER 102 **Advanced Access Control settings for Network Analysis and Intrusion Policies** 2917
Advanced access control settings for network analysis and intrusion policies 2917
Advanced access control settings for network analysis and intrusion policies 2917
Requirements and prerequisites for advanced access control settings for network analysis and intrusion policies 2918
Inspection of packets that pass before traffic is identified 2918
Best practices for handling packets that pass before traffic identification 2918
Specify a policy to handle packets that pass before traffic identification 2919
Advanced settings for network analysis policies 2920

Set the default network analysis policy 2921
Network analysis rules 2921
Network analysis policy rule conditions 2922
Configure network analysis rules 2924
Manage network analysis rules 2924

PART XXII **Objects and Certificates** 2927

CHAPTER 103 **Object Management** 2929
Objects 2930
Use the object manager 2932
Import objects 2932
Edit objects 2935
View objects and their usage 2935
Filter objects or object groups 2936
Object groups 2937
Group reusable objects 2938
Object overrides 2939
Manage object overrides 2941
Allow object overrides 2941
Add object overrides 2942
Edit an object override 2942
AAA servers 2943
Add a RADIUS server group 2943
RADIUS server group configuration options 2944
Configure RADIUS server options 2945
Add a single sign-on server 2946
Access lists 2948
Configure an extended ACL object 2949
Configure a service access object 2952
Configure a standard ACL object 2953
Configure address pools 2954
Create an application filter 2955
Configure an AS PATH object 2956

Configure a BFD template 2956
Cipher suite lists 2958
Create a cipher suite list 2958
Configure a community list 2959
Configure an extended community list 2960
DHCP IPv6 pools 2962
Distinguished names 2962
Create a distinguished name object 2964
DNS server groups 2965
Create a DNS server group object 2965
External attributes 2966
Dynamic objects 2966
Create dynamic objects with Cloud-Delivered Firewall Management Center 2966
Create Dynamic Objects with Cloud-Delivered Firewall Management Center and On-Premises Dynamic Attributes Connector 2968
Dynamic objects page features 2970
Dynamic object mappings 2970
API-created dynamic objects 2971
Add or edit an API-created dynamic object 2971
Security group tags 2972
Create a security group tag object 2972
File lists 2973
Source files for file lists 2974
Add individual SHA-256 values to file lists 2975
Upload an individual file to a file list 2975
Upload source files to file lists 2976
Edit SHA-256 values in file lists 2977
Download source files from a file list 2978
FlexConfig objects 2978
Geolocation objects 2979
Create a geolocation object 2979
Interfaces 2980
Key chains 2980
Create a key chain object 2981

Networks 2982
Network wildcard masks 2983
Create a network object 2985
Import network objects 2985
Network objects and groups 2985
PKI objects 2986
Internal certificate authority objects 2987
CA certificates and private key import 2988
Import a CA certificate and private key 2988
Generate a new CA certificate and private key 2989
Signed certificates 2989
Create an unsigned CA certificate and CSR 2990
Upload a signed certificate issued in response to a CSR 2990
CA certificates and private key downloads 2991
Download a CA certificate and private key 2992
Trusted certificate authority objects 2992
Trusted CA objects 2993
Add a trusted CA object 2993
Certificate revocation lists in trusted CA objects 2994
Add a certificate revocation list to a trusted CA object 2994
External certificate objects 2995
Add external certificate objects 2995
Internal certificate objects 2996
Add internal certificate objects 2997
Certificate enrollment objects 2997
Add certificate enrollment objects 2999
Add a certificate enrollment 3000
EST enrollment configuration options in Certificate Enrollment Object 3001
Certificate Enrollment Object SCEP options 3002
Certificate enrollment object options for ACME certificates 3003
Prerequisites for using ACME certificates 3004
ACME certificate limitations and compatibility considerations 3008
Certificate Enrollment Object certificate parameters in certificate requests 3008
Certificate Enrollment Object key options 3009

Enroll PKI certificates with weak-crypto 3010
Certificate Enrollment Object revocation options 3011
Configure a policy list 3012
Ports 3014
Create port objects 3015
Import port objects 3015
Prefix lists 3015
Configure an IPv6 prefix list 3016
Configure an IPv4 prefix list 3016
Configure a route map 3017
Security intelligence 3021
Modify security intelligence objects 3022
Global and domain security intelligence lists 3023
Security intelligence lists and multitenancy 3024
Delete entries from global security intelligence lists 3025
List and feed updates for security intelligence feeds 3025
Change the update frequency for security intelligence feeds 3026
Custom security intelligence lists and feeds 3026
List custom lists and feeds requirements 3026
List URL syntax and matching criteria for URL lists and feeds 3027
Custom security intelligence feeds 3029
Create custom security intelligence feeds 3030
Update custom security intelligence feeds manually 3030
Custom Security Intelligence lists 3031
Upload new custom security intelligence lists to the Cloud-Delivered Firewall Management Center 3032
Update custom security intelligence lists 3032
Sinkhole objects 3033
Create a sinkhole object 3033
Configure an SLA monitor 3034
Time ranges 3036
Create a time range object 3036
Time zones 3038
Tunnel zones 3038

URLs 3038
Create a URL object 3039
Variable sets 3040
Variable sets in intrusion policies 3042
Variables 3042
List predefined default variables 3043
Network variables 3045
Port variables 3046
Advanced variables 3047
Variable resets 3047
Variable sets 3048
Examples: Adding user-defined variables to default sets 3049
Examples: User-defined variables in custom sets 3050
Nesting variables 3051
Manage variable sets 3053
Create a variable set 3053
Manage variables 3054
Add variables 3055
Edit variables values in a variable set 3056
VLAN tags 3057
Create a VLAN tag object 3057
VPN 3057
Certificate map objects 3058
Secure Client custom attributes objects 3059
Configure custom attributes for Secure Client 3059
Add custom attributes to your group policy 3061
Firewall Threat Defense group policy objects 3061
Configure a group policy object 3062
Group policy general options 3063
Group policy options for Secure Client 3065
Advanced group policy option fields and values 3068
Firewall Threat Defense IPsec proposals 3069
Configure an IKEv1 IPsec proposal object 3070
Configure an IKEv2 IPsec proposal object 3071

IKE policies 3071
Configure an IKEv1 policy object 3072
Configure an IKEv2 policy object 3073
Secure client customizations 3075
File objects 3075

CHAPTER 104 Certificates 3079
Requirements and prerequisites for certificates 3079
Secure Firewall Threat Defense VPN certificate guidelines and limitations 3079
Manage Firewall Threat Defense certificates 3080
Automatically update CA bundles 3082
Install a certificate using self-signed enrollment 3084
Install a certificate using EST enrollment 3084
Install a certificate using SCEP enrollment 3085
Install a certificate using manual enrollment 3086
Install a certificate using a PKCS12 file 3087
Install a certificate using ACME enrollment 3088
Troubleshoot Threat Defense certificates 3088

PART XXIII Reference 3091

CHAPTER 105 FAQ and Support 3093
Security Cloud Control Firewall Management Platform Maintenance Schedule 3093
Navigate from Security Cloud Control to Cloud-Delivered Firewall Management Center 3094
What does the default action "Analyze all tunnel traffic" for prefiltering mean? 3094
How Security Cloud Control processes personal information 3094
Can I restore a backup from a different device? 3094
Does deploying a new prefilter policy immediately affect ongoing sessions? 3095
How do I keep my security databases and feeds up to date? 3095
What version of Secure Firewall Threat Defense can I manage with Cloud-Delivered Firewall Management Center? 3095
How do I exclude specific traffic (Webex, Zoom, etc) from the remote access VPN? 3096
How do I prevent users from accessing undesirable external network resources, such as inappropriate websites? 3097

Security Feed Questions 3097
How do I update intrusion rules (SRU/LSP)? 3097
How do I update my Cisco vulnerability database (VDB)? 3098
How do I update my Geolocation database? 3099
How do I update Security Intelligence feeds? 3099
How do I update URL reputations? 3099
How do I setup Rate-Based Attack Prevention on the FTD using Snort 2? 3100
Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI 3101
About Data Interfaces 3105
End-of-Support for management of Secure Firewall Threat Defense devices, Version 7.0.x, managed by Cloud-Delivered Firewall Management Center 3106

CHAPTER 106 **Cloud-Delivered Firewall Management Center command line reference** 3109
About the Cloud-Delivered Firewall Management Center CLI 3109
Cloud-Delivered Firewall Management Center CLI modes 3110
Cloud-Delivered Firewall Management Center CLI management commands 3110
Cloud-Delivered Firewall Management Center CLI show commands 3111
Cloud-Delivered Firewall Management Center CLI configuration commands 3112
Cloud-Delivered Firewall Management Center CLI system commands 3112

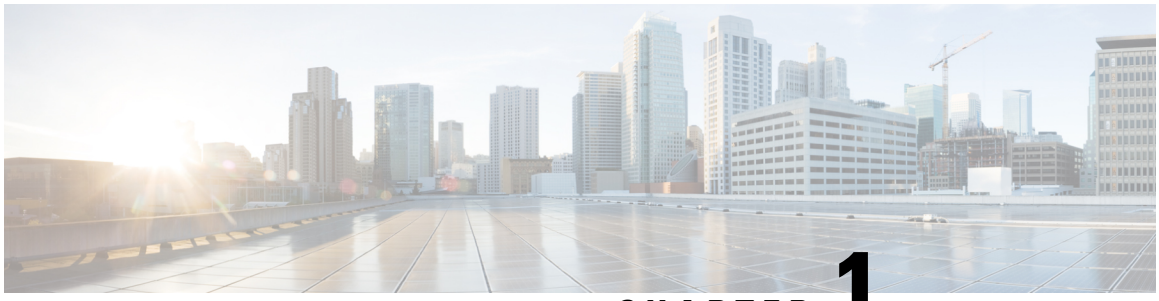
CHAPTER 107 **Communication Ports and Internet Access** 3115
Communication ports for managed devices 3115
Internet access requirements for managed devices 3116
Internet resources accessed by your browser 3118



PART I

Cisco Security Cloud Control Management: Cloud-Delivered Firewall Management Center

- Configure Cloud-Delivered Firewall Management Center-Managed Secure Firewall Threat Defense, on page 1
- Manage Multicloud Defense-Onboarded Secure Firewall Threat Defense Virtual Devices, on page 13



CHAPTER 1

Configure Cloud-Delivered Firewall Management Center-Managed Secure Firewall Threat Defense

This chapter provides information on how to manage the Cloud-Delivered Firewall Management Center-managed Firewall Threat Defense onboarded to Security Cloud Control.

- Overview of Cloud-Delivered Firewall Management Center, on page 1
- Navigate to the Cloud-Delivered Firewall Management Center in your Security Cloud Control Tenant, on page 3
- Determine the Cloud-Delivered Firewall Management Center Version in Security Cloud Control Firewall Management, on page 3
- Enable Cloud-Delivered Firewall Management Center , on page 4
- Hardware and Software Support, on page 5
- **Security Cloud Control Firewall Management Platform Maintenance Schedule**, on page 5
- Licensing, on page 6
- Integrate an On-Premises Firewall Management Center With Cisco Security Cloud, on page 6
- Security Cloud Control Firewall Management Integrations Page , on page 9

Overview of Cloud-Delivered Firewall Management Center

Cloud-Delivered Firewall Management Center is a software-as-a-service product that manages Secure Firewall Threat Defense devices through Security Cloud Control Firewall Management. It provides many of the same functions as an On-Premises Firewall Management Center.

Cloud-Delivered Firewall Management Center has the same appearance and behavior as an On-Premises Firewall Management Center and uses the same FMC API.

Because it is a SaaS product, the Security Cloud Control Firewall Management operations team deploys and maintains the Cloud-Delivered Firewall Management Center software. As new features are introduced, the operations team updates the Cloud-Delivered Firewall Management Center in your Security Cloud Control Firewall Management organization.

Migration and onboarding details

A migration wizard is available to help you migrate Threat Defense devices from an On-Premises Firewall Management Center to Cloud-Delivered Firewall Management Center.

Devices must run one of these Threat Defense software releases before you migrate them:

- Minimum On-Premises Firewall Management Center: 7.2
- Minimum Threat Defense: 7.2.x (not supported for Version 7.1)

Threat Defense 7.1 releases are not supported for migration.

You onboard Threat Defense devices in Security Cloud Control Firewall Management by using familiar methods, such as entering the serial number or using a CLI command that includes a registration key.

After onboarding:

- The device appears in both Security Cloud Control Firewall Management and Cloud-Delivered Firewall Management Center.
- You configure the device in Cloud-Delivered Firewall Management Center.
- In Security Cloud Control Firewall Management, you can review device information such as version, configuration status, connectivity, health status, and node status.
- When you click the health status in Security Cloud Control Firewall Management, the platform opens the device health monitoring page in the Cloud-Delivered Firewall Management Center user interface.

Availability and Security Analytics

Security Cloud Control Firewall Management supports high availability for the Threat Defense devices that it manages through the data interface. This feature is supported for devices that run software version 7.2 or later.

You can analyze security events from onboarded threat defense devices by using either Security Analytics and Logging (SaaS) or Security Analytics and Logging (On-Premises):

- Security Analytics and Logging (SaaS) stores events in the cloud, and you view those events in Security Cloud Control Firewall Management.
- Security Analytics and Logging (On-Premises) stores events on an on-premises Secure Network Analytics appliance, and analysis occurs in the On-Premises Firewall Management Center.

In both cases, you can still send logs directly from firewalls to a log collector of your choice.

Licenses

The license for Cloud-Delivered Firewall Management Center is a per-device-managed license and there is no license required for the Cloud-Delivered Firewall Management Center itself. Existing Secure Firewall Threat Defense devices re-use their existing smart licenses and new Secure Firewall Threat Defense devices provision new smart licenses for each feature implemented on the Secure Firewall Threat Defense For more information, see Cloud-Delivered Firewall Management Center license.

The **Unified Events** feature in Cloud-Delivered Firewall Management Center uses Cisco Security Analytics and Logging (SaaS) as its event data source. You must have a valid Cisco Security Analytics and Logging (SaaS) subscription plan to view firewall events on the **Unified Events** page. When you enable Cloud-Delivered Firewall Management Center, a 90-day trial subscription for Security Analytics and Logging (SaaS) is automatically issued.

To learn how to have a Cloud-Delivered Firewall Management Center provisioned on your tenant, see Enable Cloud-delivered Firewall Management Center on Your Security Cloud Control Tenant.

Navigate to the Cloud-Delivered Firewall Management Center in your Security Cloud Control Tenant

Procedure

Step 1 In the left pane, click **Administration** > **Integrations** > **Firewall Management Center** to view the **Services** page.

The **Services** page displays the information about the Cloud-Delivered Firewall Management Center provisioned for the Security Cloud Control tenant.

Step 2 Choose **Cloud-Delivered FMC** and click the links in the **Actions**, **Management**, or **Settings** pane to navigate to Cloud-Delivered Firewall Management Center to perform various actions. See Security Cloud Control Firewall Management Integrations Page , on page 9.

Once in the Cloud-Delivered Firewall Management Center, click the blue question mark button and select **Page-level Help** to learn more about the page you are on and further action you can take.

Determine the Cloud-Delivered Firewall Management Center Version in Security Cloud Control Firewall Management

Information about the Cloud-Delivered Firewall Management Center version is required primarily for migration planning, feature support verification, troubleshooting, licensing, and operational management.

Use this procedure to determine the current version of your Cloud-Delivered Firewall Management Center in Security Cloud Control Firewall Management.

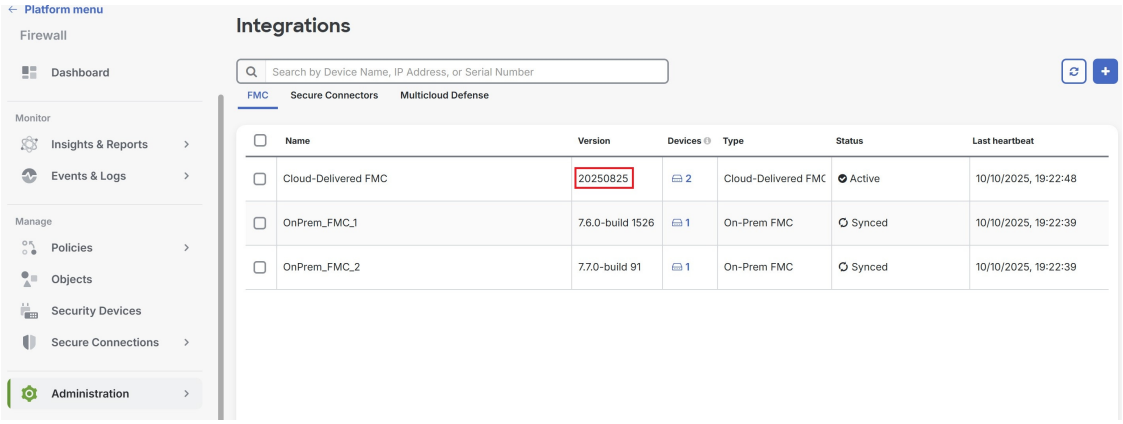
Before you begin

Enable Enable Cloud-Delivered Firewall Management Center in your organzation.

Procedure

Step 1 In the left pane, click **Administration** > **Integrations** > **Firewall Management Center**.

Step 2 Under the **FMC** tab, the version is displayed under the **Version** column.



For more information about Cloud-Delivered Firewall Management Center versions, see Release Notes for Cloud-Delivered Firewall Management Center.

Enable Cloud-Delivered Firewall Management Center

If you want to manage your Secure Firewall Threat Defense devices, you can enable the Cloud-Delivered Firewall Management Center on your tenant. You need to have an admin or a super admin user role to perform this task.

Procedure

From the Security Cloud Control menu, click **Administration** > **Integrations** > **Firewall Management Center** and click **Enable Cloud-Delivered FMC**.

Security Cloud Control starts provisioning a Cloud-Delivered Firewall Management Center instance in the background; it typically takes 15 to 30 minutes for this to be complete. You can track the provisioning progress on the **Status** column of **Cloud-Delivered FMC**.

After the provisioning is complete, the status changes to **Active**. In addition, you get a **Cloud-Delivered Firewall Management Center is Ready** notification on the Security Cloud Control notifications panel and on the applications on which you have configured incoming webhooks. See Notification Settings for more information.

Note
After you receive the **Cloud-Delivered Firewall Management Center is Ready** notification, ensure that you log out of and log in back to your tenant once, to see the **Cloud-Delivered FMC** right pane options, such as **Actions**, **Management**, and **System**.

What to do next:

a. Register Cloud-Delivered Firewall Management Center with Cisco Smart Software Manager (CSSM) to activate smart license. For detailed steps, refer to Register the Cloud-Delivered Firewall Management Center with the Smart Software Manager.

b. Onboard your Firewall Threat Defense devices to the Cloud-Delivered Firewall Management Center and manage them.

Hardware and Software Support

Cloud-Delivered Firewall Management Center supports these Secure Firewall Threat Defense software versions when they are installed on any supported hardware or virtual device:

- Verion 7.0.3 or later 7.0.x versions.
- Version 7.2 and later versions.

Note Software Version 7.1 is _not_ supported.

See Firepower Threat Defense Support Specifics for more information.

Security Cloud Control Firewall Management Platform Maintenance Schedule

Security Cloud Control Firewall Management updates its platform every week with new features and quality improvements. Updates are made during a 3-hour period according to this schedule:

Day of the week	Time of day (24-hour time, UTC)
Thursday	09:00 UTC - 12:00 UTC

During the Security Cloud Control Firewall Management upgrade period, you can still access your organization and the Cloud-Delivered Firewall Management Center. Additionally, the devices that you have onboarded to Security Cloud Control Firewall Management continue to enforce their security policies.

Note

- During the maintenance period of the Security Cloud Control Firewall Management, refrain from creating or deploying configurations to the devices it manages.
- If there is any issue that stops Security Cloud Control Firewall Management from communicating, Cisco addresses that issue in all the affected tenants as quickly as possible, even if it is outside the maintenance window.

Licensing

The **Unified Events** page uses the Cisco Security Analytics and Logging data store as its event data source. You must have a valid Cisco Security Analytics and Logging subscription plan to view firewall events on the **Unified Events** page.

Integrate an On-Premises Firewall Management Center With Cisco Security Cloud

This procedure describes how to integrate the On-Premises Firewall Management Center with Cisco Security Cloud. By enabling Cisco Security Cloud integration, your management center gets registered to the Cisco cloud tenancy.

Before you begin

When onboarding an on-premises Firewall Management Center:

- If you already have a Security Cloud Control tenant, you will be prompted to confirm whether you want to attach the on-premises Firewall Management Center to that existing tenant.
- If you do not have a tenant or account, the on-premises Firewall Management Center integration wizard will automatically create a new tenant for you during the registration task the onboarding process. This tenant is a free tier provided exclusively for on-premises Firewall Management Center AI Assist enablement. Allow the wizard to manage tenant creation so that you do not receive a trial tenant that would later require conversion. The wizard manages this process to provide a seamless experience.
- If you intend to use the Cloud-Delivered Firewall Management Center or manage firewalls using the Security Cloud Control, you must purchase a base license and device entitlements. For detailed licensing information and product options, refer to the Security Cloud Control Licensing Overview.

Procedure

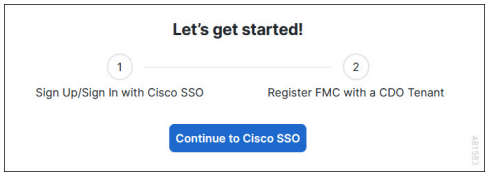
Step 1 Based on your on-premises Firewall Management Center version, do the following:

- For on-premises Firewall Management Center version between 7.2 and 7.4.x, click **Enable Secure X**.
- For on-premises Firewall Management Center version 7.6 or later, click **Cisco Security Cloud**.
- For on-premises Firewall Management Center version 10.0 or later, click **Security Cloud Control**.

A separate browser tab opens to log you in to your Security Cloud Control account. Make sure this page is not blocked by a pop-up blocker.

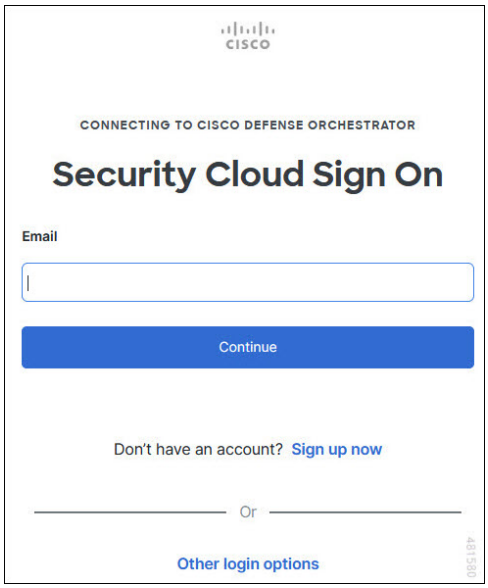
Step 2 Click **Continue to Cisco SSO**.

Figure 1: Cisco Security Cloud Welcome Page



Step 3 Log in to your Security Cloud Control account.

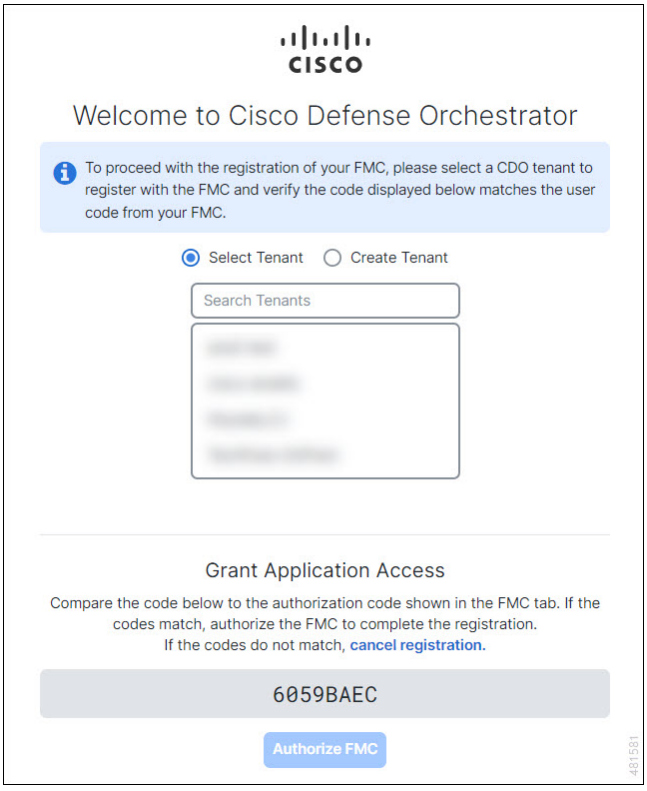
Figure 2: Cisco Security Cloud Sign On



If you do not have a security cloud sign on account to log in to Security Cloud Control and you want to create one, click **Sign up now** in the **Security Cloud Sign On** page.

Step 4 Choose a Security Cloud Control tenant that you want to use for this integration. The on-premises Firewall Management Center and the managed devices get onboarded to the Security Cloud Control tenant that you choose here.

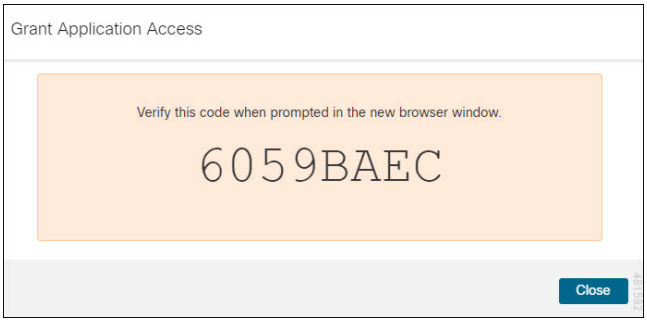
Figure 3: Choose the Security Cloud Control Tenant



If you do not have a security cloud sign on account to log in to Security Cloud Control and you want to create one, click **Sign up now** in the **Security Cloud Sign On** page. See Create a New Cisco Security Cloud Sign On Account.

Step 5 Verify that the code displayed in the Security Cloud Control login page matches the code provided by the on-premises Firewall Management Center.

Figure 4: Verification Code in the on-premises Firewall Management Center



Step 6 Click **Authorize FMC**.

Step 7 In the on-premises Firewall Management Center UI, click **Save** to save the configuration.

You can view the task progress under **Notifications** > **Tasks**.

The registration task can take up to 90 seconds to complete. If you must use on-premises Firewall Management Center while the registration task is in progress, open the on-premises Firewall Management Center in a new window.

Security Cloud Control Firewall Management Integrations Page

The **Integrations** page provides a comprehensive view of your managed infrastructure within Security Cloud Control Firewall Management:

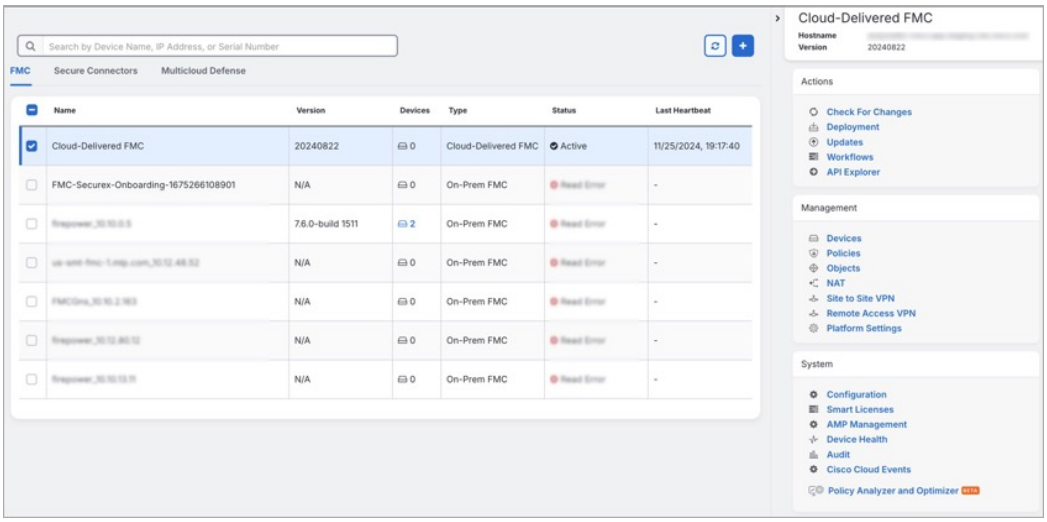
- **FMC** tab:
  - View the Cloud-Delivered Firewall Management Center that is linked to the Security Cloud Control Firewall Management account. Refer to Request a Cloud-delivered Firewall Management Center for your Security Cloud Control tenant .
  - View which on-premises Firewall Management Centers are onboarded to Security Cloud Control Firewall Management. Devices managed by these on-premises Firewall Management Centers appear in **Security Devices**. Refer to Managing On-Prem Firewall Management Center with Security Cloud Control.
  - Lists the **Secure Device Connector** and **Secure Event Connector** available under the **Secure Connectors**. Refer to Secure Device Connector and Secure Event Connectors.

Security Cloud Control. page. The **Integrations** page also lists tab.

Click the **FMC** tab, then click **Onboard FMC** () icon to onboard an on-premises Firewall Management Center. In the right pane, perform device actions as needed. Refer to Onboard an On-Prem Firewall Management Center.

View device information, including version, device count, device type, and synchronization status. Click the managed devices icon to go to the **Security Devices** page. Devices managed by the selected on-premises Firewall Management Center are filtered and displayed. On the **Integrations** page, select multiple on-premises Firewall Management Center to perform actions on several management centers at once. Deselect the Cloud-Delivered Firewall Management Center before selecting other on-premises Firewall Management Center. To add a new secure connector or perform actions on existing secure connectors, choose the **Secure Connectors** tab and click .

In the left pane, click **Administration** > **Firewall Management Center**.



For your Cloud-Delivered Firewall Management Center, the Integrations page displays this information:

- If you do not have a Cloud-Delivered Firewall Management Center deployed on your tenant, click **Enable Cloud-Delivered FMC**. Refer to Enable Cloud-Delivered Firewall Management Center on Your Security Cloud Control Tenant for more information.
- The number of Secure Firewall Threat Defense devices deployed on the Cloud-Delivered Firewall Management Center.
- Status of the connection between Security Cloud Control and the Cloud-Delivered Firewall Management Center page.
- The last heartbeat of the Cloud-Delivered Firewall Management Center. This time shows when the Cloud-Delivered Firewall Management Center's status and device count last synchronized with the table on this page..
- The hostname of the selected Cloud-Delivered Firewall Management Center.

Choose **Cloud-Delivered FMC** and using the links in the **Actions** , **Management**, or **Settings** pane, you open the Cloud-Delivered Firewall Management Center user interface to perform the configuration tasks that are associated with the link you clicked.

Actions:

- **Check For Changes**: The Device Count and Status information in the table is updated with the information available the last time this page and the Cloud-Delivered Firewall Management Center were synchronized. Synchronization happens every 10 minutes.
- **Deployment**: Takes you to the device configuration deployment page on Cloud-Delivered Firewall Management Center. Refer to Deploy Configuration Changes.
- **Workflows**: Takes you to the **Workflows** page to monitor every process that Security Cloud Control runs when communicating with devices. Refer to Workflows page.
- **API Explorer**: Takes you to the page that lists the Cloud-Delivered Firewall Management Center REST APIs. Refer to Secure Firewall Management Center REST API Guide .
- **Unified Events**: Takes you to the **Unified Events** page on the Cloud-Delivered Firewall Management Center portal, which provides a single-screen view of various firewall events, including connection,

intrusion, file, malware, and security-related connection events. For more information, refer to Unified Events .

Note The Unified Events feature requires activation. If you have not yet activated this feature, contact your Cisco sales representative to enable it.

Management:

- **Devices**: Takes you to the Firewall Threat Defense device listing page on the Cloud-Delivered Firewall Management Center portal. Refer to Configure Devices.
- **Policies**: Takes you to the policies page on the Cloud-Delivered Firewall Management Center portal to edit system-provided access control policies and create custom access control policies. Refer to Manage Access Control Policies.
- **Objects**: Takes you to the policies page on the Cloud-Delivered Firewall Management Center portal to manage reusable objects. Refer to Object Management.
- **NAT**: Takes you to the policies page on the Cloud-Delivered Firewall Management Center portal to configure Network Address Translation policies on the Firewall Threat Defense devices. Refer to Manage NAT policies.
- **Site to Site VPN**: Takes you to the site-to-site VPN dashboard page on the Cloud-Delivered Firewall Management Center portal to configure site-to-site VPN policy between two sites. Refer to Site-to-Site VPNs.
- **Remote Access VPN**: Takes you to the remote access VPN dashboard page on the Cloud-Delivered Firewall Management Center portal to configure a remote access VPN configuration. Refer to Remote Access VPN.
- **Platform Settings**: Takes you to the platform settings page on the Cloud-Delivered Firewall Management Center portal to configure a range of unrelated features whose values you might want to share among several devices. Refer to Platform Settings.

System:

- **Configuration**: Takes you to the system configuration settings page on the Cloud-Delivered Firewall Management Center portal to configure system configuration settings. Refer to System Configuration .
- **Smart Licenses**: Takes you to the smart licenses page on the Cloud-Delivered Firewall Management Center portal to assign licenses to devices. Refer to Assign Licenses to Devices.
- **AMP Management**: Takes you to the AMP management page on the Cloud-Delivered Firewall Management Center portal that supplies data the system uses to detect and block malware on your network. Refer to Cloud Connections for Malware Protection.
- **Device Health**: Takes you to the health monitoring page on the Cloud-Delivered Firewall Management Center portal, which tracks various health indicators to ensure the hardware and software in the system are working correctly. Refer to About Health Monitoring.
- **Audit**: Takes you to the audit log page on the Cloud-Delivered Firewall Management Center portal to show the generated audit record for each user interaction with the web interface.

- **Cisco Cloud Events**: Takes you to the configure Cisco Cloud events page on the Security Cloud Control portal to configure Cloud-Delivered Firewall Management Center to send events directly to SAL (SaaS). Refer to Send Events to SAL (SaaS).

After you open the Cloud-Delivered Firewall Management Center page, click the blue question mark button and select **Page-level Help** for more information.

Support to open Security Cloud Control and Cloud-Delivered Firewall Management Center applications in separate tabs

As you configure Firewall Threat Defense devices or objects in Cloud-Delivered Firewall Management Center , you can open the appropriate configuration pages in additional browser tabs to work simultaneously in the Security Cloud Control and the Cloud-Delivered Firewall Management Center portals without logging off.

For example, you can create an object on Cloud-Delivered Firewall Management Center and simultaneously monitor event logs on Security Cloud Control that are generated from the security policies.

This feature is available for all Security Cloud Control links that navigate to the Cloud-Delivered Firewall Management Center portal. To open the Cloud-Delivered Firewall Management Center portal in a new tab:

On the Security Cloud Control portal, press and hold the Ctrl (Windows) or Command (Mac) button, then click the corresponding link.

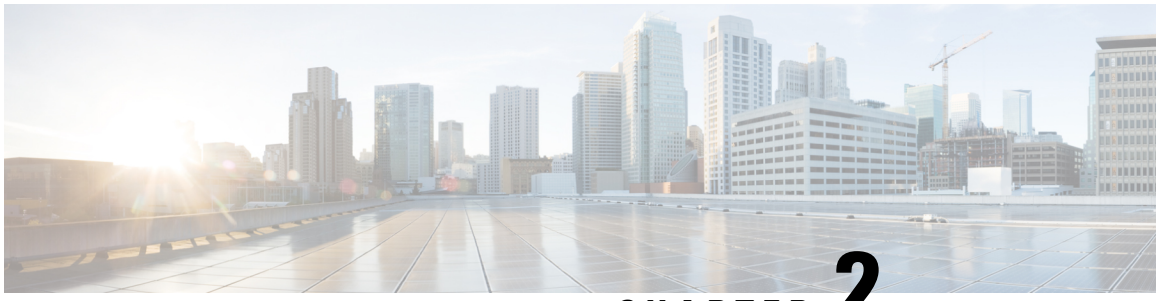
Note A single click opens the Cloud-Delivered Firewall Management Center page in the same tab.

Here are some examples of opening the Cloud-Delivered Firewall Management Center portal page in a new tab:

- Choose **Administration** > **Firewall Management Center** and select **Cloud-Delivered FMC**. In the right pane, press and hold the Ctrl (Windows) or Command (Mac) button, and then click the page that you want to access.
- Choose **Objects** > **Other FTD Objects**.
- Click the search icon on the Security Cloud Control page and enter your search terms in the search field.

  From the search result, press and hold the Ctrl (Windows) or Command (Mac) button, and then click the arrow icon.
- Choose **Dashboard** > **Quick Actions**. Press and hold the Ctrl (Windows) or Command (Mac) button, and then click **Manage FTD Policies** or **Manage FTD Objects**.

Note If you switch to a new Security Cloud Control tenant, the Cloud-Delivered Firewall Management Center portal open in another tab logs out automatically.



CHAPTER 2

Manage Multicloud Defense-Onboarded Secure Firewall Threat Defense Virtual Devices

- Overview of Multicloud Defense-Onboarded Firewall Threat Defense Virtual Devices, on page 13
- Onboard and Configure a Secure Firewall Threat Defense Virtual Device in Multicloud Defense, on page 15

Overview of Multicloud Defense-Onboarded Firewall Threat Defense Virtual Devices

You can orchestrate your Secure Firewall Threat Defense virtual (FTDv) device in your cloud service provider manually and then register your device to Cloud-Delivered Firewall Management Center (Cloud-Delivered Firewall Management Center) for management of it. A lot of manual steps are involved in this method such as VPC creation, security groups creation, peer connection and routing in your cloud service provider, which can be cumbersome.

With Multicloud Defense available as part of Security Cloud Control, you can now easily orchestrate, deploy, and register an FTDv device using Multicloud Defense while retaining your FTDv as a gateway. Multicloud Defense communicates with Security Cloud Control for FTDv registration and orchestrates the initial setting up of the FTDv. For more details on deploying an FTDv using Multicloud Defense, see Secure Firewall Threat Defense.

HTTPS configuration under Platform Settings policy is managed by Multicloud Defense. Any other configurations related to Platform Settings need to be done in Firewall Management Center only. Multicloud Defense also adds one mandatory NAT rule under NAT policy. Additional NAT rules can be configured in the Firewall Management Center. You can create and assign access policies, rules, and objects in Cloud-Delivered Firewall Management Center For more details, see Access Control Policies.

Guidelines for Managing an FTDv Created in Multicloud Defense

Consider the following guidelines in the context of Cloud-Delivered Firewall Management Center, when creating gateways for FTDv using Multicloud Defense:

- FTDv registration and configuration can be done only in Multicloud Defense. This information is view-only and cannot be edited in Cloud-Delivered Firewall Management Center. You cannot edit any of the **Devices**, **Interface**, **Inline Sets**, **Routing**, **DHCP** and **VTEP** configuration details in the user interface of Cloud-Delivered Firewall Management Center, because all device-specific configurations are done in Multicloud Defense.

- HTTPS configuration under Platform Settings policy is managed by Multicloud Defense and any changes to it can impact devices or the FTDv gateway. You need to be cautious while changing any configurations in **Platform Settings** as the policy overall is a shared policy.
- Mandatory NAT rule is managed by Multicloud Defense and any changes to the mandatory rule can impact devices or the FTDv gateway.
- Configure NAT rules to translate incoming traffic from the external network to the internal network on the specified port and IP address. Multicloud Defense automatically opens the required port on the load balancer provided by the cloud service. This configuration ensures accurate translation and forwarding of inbound traffic, while the load balancer manages port access.
- Multicloud Defense devices are added under a logical device group for each gateway.
- You can use existing policies that you have in Cloud-Delivered Firewall Management Center or even create new policies from the Multicloud Defense application.
- You can configure or modify an existing policy from Cloud-Delivered Firewall Management Center.
- To add additional policies such as QoS policy or any other policy configuration for Multicloud Defense-related FTDv in Cloud-Delivered Firewall Management Center, the policies must be assigned at the device group level, so that it is applied to all devices in the group. This helps in addressing shared policies and the application of the policy to all devices that are added to the group in future. Note that policies are not applied at the single FTDv device level, but across the group.
- Don't remove an FTDv from Security Cloud Control. Contact your administrator to make any changes, if any, in Multicloud Defense.
- Don't move an FTDv across groups at any point in time or unassign an FTDv from groups in Cloud-Delivered Firewall Management Center.
- Don’t rename the device groups created by Multicloud Defense in Cloud-Delivered Firewall Management Center, because it can hamper the working of the devices.
- Don’t edit or delete the security zones or objects because they are managed in Multicloud Defense and can hamper the working of the devices. You can identify this when you access Platform Settings policies. If the policy is deployed to a device orchestrated by Multicloud Defense, you are notified with a message. For identification, Cloud-Delivered Firewall Management Center uses an additional field with the naming convention `mcdmanaged`. You can identify the security zones by the naming convention used – `ciscomcd-VNI`, `ciscomcd-inside`, or `ciscomcd-outside`.
- Licensing of the FTDv is according to the FTDv plan that you chose, that is, either Bring Your Own License (BYOL) or Software-as-a-Service (SaaS). You can use the existing license that you use for FTDv.
- We recommend that you upgrade the FTDv versions only in the Multicloud Defense application.
- This feature is supported for FTDv version 7.6 and above. You can upgrade or downgrade the FTDv in Multicloud Defense. This is done by creating a new instance. The upgrade or downgrade enables easy scaling up or scaling down of the instances too.
- Your FTDv is listed in Security Cloud Control Firewall Management and Cloud-Delivered Firewall Management Center. You can manage your policies in Cloud-Delivered Firewall Management Center.
- When you trigger a deployment on these FTDvs, you should select together all the FTDvs that belong to same gateway, otherwise the configuration will be out of sync among FTDvs and this can lead to inconsistent traffic behavior.

- Features such as Site-to-Site VPN and Remote Access VPN are not currently supported.
- Multicloud Defense can register FTDv devices with Cloud-Delivered Firewall Management Center by using the supported registration flow for the deployment. For Multicloud Defense-managed FTDv devices, registration can be device-initiated or Cloud-Delivered Firewall Management Center-triggered. In the Cloud-Delivered Firewall Management Center-triggered flow, Cloud-Delivered Firewall Management Center initiates the registration connection to the FTDv by using the registration information that Multicloud Defense provides through Security Cloud Control Firewall Management.

  This support applies only to FTDv devices that are deployed and managed by Multicloud Defense. It does not change the onboarding flow for manually onboarded FTD or FTDv devices in Cloud-Delivered Firewall Management Center.

Onboard and Configure a Secure Firewall Threat Defense Virtual Device in Multicloud Defense

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 In the Security Cloud Control platform menu, choose **Products > Multicloud Defense.**

Step 3 Register your FTDv in Multicloud Defense and onboard your cloud security provider, either AWS or Azure. For more details, see Secure Firewall Threat Defense in the Cisco Multicloud Defense User Guide.

Step 4 Create a Service VPC or VNet in Multicloud Defense. For more details, see Create a VPC or VNet.

Step 5 Configure your gateway in Multicloud Defense. For more details, see Add a Gateway with an FTDv.

Step 6 Configure your policy in Cloud-Delivered Firewall Management Center. For more details, see Access Control Policies.

What to do next

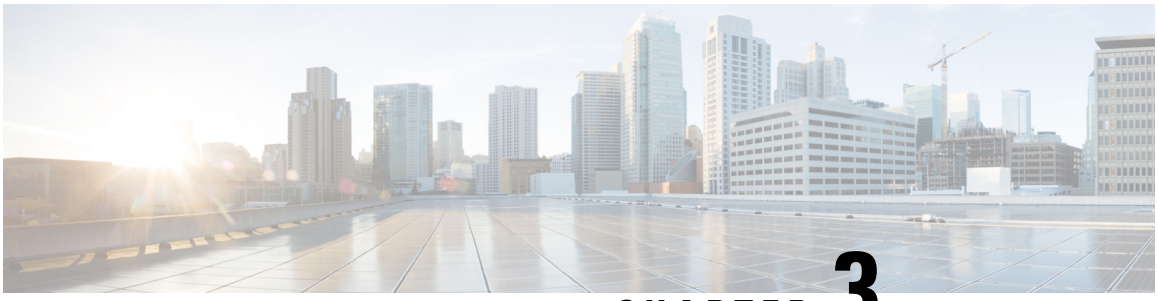
You can view your FTDv in Multicloud Defense, Security Cloud Control platform, and Firewall Management Center in your Cloud-Delivered Firewall Management Center account.



PART II

Onboard Devices to Cloud-Delivered Firewall Management Center

- Onboard a Secure Firewall Threat Defense to the Cloud-Delivered Firewall Management Center, on page 19
- Migrate On-Premises Firewall Management Center Managed Secure Firewall Threat Defense to Cloud-delivered Firewall Management Center, on page 63
- Device Registration, on page 99
- Device Registration Using Device Templates, on page 119
- Transparent or Routed Firewall Mode, on page 153
- Change Management, on page 167
- CLI Users, on page 181
- Configuration Deployment, on page 189



CHAPTER 3

Onboard a Secure Firewall Threat Defense to the Cloud-Delivered Firewall Management Center

Read the following information for onboarding prerequisties and procedures.

- Onboarding Overview, on page 19
- Maximum Firewall Threat Defense devices supported in Cloud-Delivered Firewall Management Center, on page 20
- Prerequisites to Onboard a Device to Cloud-Delivered Firewall Management Center, on page 21
- Onboard a Device with a CLI Registration Key, on page 22
- Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning, on page 24
- Onboard a Firewall Threat Defense Device to On-Prem Firewall Management Center using Zero-Touch Provisioning, on page 25
- Onboard Firewall Threat Defense Devices using Device Templates to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning, on page 27
- Cloud Deployment, on page 29
- Onboard a Secure Firewall Threat Defense Cluster, on page 46
- Onboard a Chassis, on page 49
- Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI, on page 51
- Delete Devices from Cloud-Delivered Firewall Management Center, on page 54
- Troubleshooting, on page 55

Onboarding Overview

Review the following use cases and supported sofware versions that are compatible with Cloud-Delivered Firewall Management Center management.

Note You must ensure that the Firewall Threat Defense device ports have external and outbound access for the Cloud-Delivered Firewall Management Center to onboard them. There is no requirement for an on-premises or cloud-based Security Device Controller (SDC) for this operation. For more information, see Network Requirements.

To send Firewall Threat Defense Syslog events to the Cisco cloud, you can set up the Secure Event Connector (SEC). For more information, see Installing Secure Event Connectors.

Firewall Threat Defense Devices Currently Managed by Cloud-Delivered Firewall Management Center

The following scenarios occur when you either move or migrate a device to the Cloud-Delivered Firewall Management Center:

- If you delete a device from an on-premises Firewall Management Center or Secure Firewall Threat Defense Firewall Device Manager to onboard to the Cloud-Delivered Firewall Management Center, the change of managers wipes any policies configured through the on-premises Firewall Management Center.
- If you **migrate** a device from an on-premises Firewall Management Center to the Cloud-Delivered Firewall Management Center, the device retains the majority of your previously configured policies.

Note If you do not know if your device is already managed by an alternative manager, use the `show managers` command in the device's CLI.

Onboarding Methods

Cloud-Delivered Firewall Management Center supports the following onboarding methods:

- Registration Key - Onboard a device with a registration key. The initial device setup wizard is complete on the device.
- Zero-Touch Provisioning - Onboard a new factory-shipped device with its serial number. Note that this method only supports Firepower 1000, Firepower 2100, or Secure Firewall 3100 devices.
- Zero-Touch Provisioning using a device template - Onboard new factory-shipped devices using serial numbers and a device template. Note that this method only supports Firepower 1000, Firepower 2100, Secure Firewall 1200 or Secure Firewall 3100 devices.

Maximum Firewall Threat Defense devices supported in Cloud-Delivered Firewall Management Center

Cloud-Delivered Firewall Management Center supports up to 2,000 Firewall Threat Defense devices per deployment. This total can include a combination of standalone devices, devices in high-availability pairs, and devices that are cluster nodes.

The Cloud-Delivered Firewall Management Center can manage Firewall Threat Defense Version 7.2.0 and later.

Note The maximum number of supported Secure Firewall Threat Defense devices in Cloud-Delivered Firewall Management Center includes devices that are deployed for analytics on Firewall Management Center and onboarded to Cloud-Delivered Firewall Management Center for management. These devices count toward the platform device capacity limit and must be included in your capacity planning to ensure proper sizing and performance.

Prerequisites to Onboard a Device to Cloud-Delivered Firewall Management Center

Onboard Limitations and Requirements

Be aware of the following limitations when onboarding a device to the Cloud-Delivered Firewall Management Center:

- Devices **must** be running version 7.2 and later.
- You can migrate an HA pair that is managed by an On-Premises Firewall Management Center by following the Migrate FTD to Cloud-Delivered Firewall Management Center process. Confirm both peers are in a healthy state prior to migrating.
- Devices managed locally by Firewall Device Manager can be onboarded using the serial number.
- Firewall Threat Defense devices can be onboarded to Cloud-Delivered Firewall Management Center using zero-touch provisioning if they are new (just unboxed) or have undergone a factory reset.
- If the device is managed by an on-premises Firewall Management Center, you can either onboard the device to Cloud-Delivered Firewall Management Center or migrate the device. Migrating retains any existing policies and objects, whereas onboarding the device removes most policies and all objects. See Migrate FTD to Cloud-Delivered Firewall Management Center for more information.
- If your device is currently managed by a Firewall Device Manager, unregister all your smart licenses before you onboard the device. Even if you switch device management, the Cisco Smart Software Manager will retain the smart licenses.
- If you have previously onboarded a device that was managed by a Firewall Device Manager and deleted the device from Security Cloud Control with the intention of re-onboarding for cloud management, you **must** register the Firewall Device Manager to the Security Services Exchange cloud after deleting the device. See the "Access Security Services Exchange" chapter in the *Firepower and Cisco SecureX Threat Response Integration Guide.*

Tip Onboarding a device to the Cloud-Delivered Firewall Management Center removes any policies and most objects configured through the previous manager. If your device is currently managed by an on-premises Firewall Management Center, it is possible to migrate the device and retain your policies and objects. See Migrate FTD to Cloud-Delivered Firewall Management Center for more information.

Network Requirements

Before you onboard a device, ensure the following ports have external and outbound access. Confirm the following ports on the device are allowed. If communication ports are blocked behind a firewall, onboarding the device may fail.



Note You cannot configure these ports in the Security Cloud Control UI. You must enable these ports through the device's SSH.

Table 1: Device Port Requirements

Port	Protocol / Feature	Details
443/tcp	HTTPS	Send and receive data from the internet.
443	HTTPS	Communicate with the AMP cloud (public or private)
8305/tcp	Appliance communications	Securely communicate between appliances in a deployment.

Management and Data Interfaces

Make sure your device is correctly configured with either a management or data interface.

To configure a management or data interface on your device, see Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI, on page 51.

Onboard a Device with a CLI Registration Key

Use the procedure below to onboard a device for Cloud-Delivered Firewall Management Center with a CLI registration key.



Note If your device is currently managed by an on-premises Firewall Management Center, onboarding the device will fail. You can either delete the device from the on-premises Firewall Management Center and onboard as a fresh, new device with no policies or objects, or you can migrate the device and retain the existing policies and objects. See Migrate FTD to Cloud-Delivered Firewall Managmenet Center for more information.



Important You can create a Security Cloud Control-managed, standalone logical Firewall Threat Defense device using the Secure Firewall chassis manager or the FXOS CLI.

Before you begin

Before you onboard a device, be sure to complete the following tasks:

- Cloud-Delivered Firewall Management Center is enabled for your tenant.
- Confirm the device's CLI configuration is successfully completed. See Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI, on page 51 for more information.
- Review the prerequesites and limitations before you onboard the device. See "Prerequisites to Onboard a Device to Cloud-Delivered Firewall Management Center" in Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control for more information.
- The device can be configured for either local management with Secure Firewall Device Manager or remote management with Cloud-Delivered Firewall Management Center.
- Device must be running version 7.2.0 and later.
- You have reset the device's SSH password as part of the bootstrap process. If you have you not reset the SSH password, Security Cloud Control recommends using the Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning, on page 24 method.

Procedure

Step 1 In the left pane, click **Security Devices**.

Step 2 Click **Onboard device or service** () icon.

Step 3 Click the **FTD** tile.

Step 4 Under **Management Mode**, select **FTD**. By selecting **FTD** under **Management Mode**, you will not be able to manage the device using the previous management platform. All existing policy configurations except for interface configurations will be reset. You must re-configure policies after you onboard the device.

Step 5 Select **Use CLI Registration Key** as the onboarding method.

Step 6 Enter the device name in the **Device Name** field and click **Next**.

Step 7 In the **Policy Assignment** step, use the drop-down menu to select an access control policy to deploy once the device is onboarded. If you have no policies configured, select the **Default Access Control Policy**.

Step 8 Specify whether the device you are onboarding is a physical or virtual device. If you are onboarding a virtual device, you must select the device's performance tier from the drop-down menu.

Step 9 Select the subscription licenses you want to apply to the device. Click **Next**.

Step 10 Security Cloud Control generates a command with the registration key. Connect to the device you are onboarding using SSH. Log in as "admin" or a user with equivalent admin privileges and paste the entire registration key as is into the device's CLI.

Note: For Firepower 1000, Firepower 2100, ISA 3000, and Firewall Threat Defense Virtual devices, open an SSH connection to the device and log in as `admin`. Copy the entire registration command and paste it into the device's CLI interface at the prompt. In the CLI, enter **Y** to complete the registration. If your device was previously managed by Firewall Device Manager, enter **Yes** to confirm the submission.

Step 11 Click **Next** in the Security Cloud Control onboarding wizard.

Step 12 (Optional) Add labels to your device to help sort and filter the **Security Devices** page. Enter a label and select the blue plus button. Labels are applied to the device after it's onboarded to Security Cloud Control.

Once the device is synchronized, select the device you just onboarded from the **Security Devices** page and select any of the options listed under the **Device Management** pane located to the right. We strongly recommend the following actions:

- If you did not already, create a custom access control policy to customize the security for your environment. See Access Control Overview in *Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control* for more information.
- Enable Cisco Security Analytics and Logging (SAL) to view events in the Security Cloud Control dashboard **or** register the device to an Secure Firewall Management Center for security analytics. See Cisco Security Analytics and Logging in *Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control* for more information.

Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning

Only the Firepower 1000, Firepower 2100, Secure Firewall 1200, and Secure Firewall 3100 devices can be onboarded with the zero-touch provisioning method.

Before you begin

Confirm that the following is completed before onboarding:

- You have a Security Cloud Control tenant. If you do not, see Request a Security Cloud Control Tenant for more information.
- Cloud-Delivered Firewall Management Center is enabled for your tenant.
- The device is freshly installed but has never been logged into by either the device CLI, a Firewall Management Center, or the Firewall Device Manager.
- The device is running version 7.2 or later.
- The device's outside or management interface should be able to access the Security Services Exchange domain registration.us.sse.itd.cisco.com.

Procedure

Step 1 In the left pane, click **Security Devices**.

Step 2 Click **Onboard device or service** () icon.

Step 3 Click the **FTD** tile.

Step 4 Under **Management Mode**, select **FTD**. By selecting **FTD** under **Management Mode**, you will not be able to manage the device using the previous management platform. All existing policy configurations except for interface configurations will be reset. You must re-configure policies after you onboard the device.

Step 5 Click the **Use Serial Number** tile.

Step 6 Select Cloud-Delivered Firewall Management Center from the drop-down list. Click **Next**.

For information on onboarding a threat defense device to On-Prem Firewall Management Center, see Onboard a Threat Defense Device to On-Prem Firewall Management Center using Zero-Touch Provisioning.

Step 7 Enter the **Device Serial Number** and the **Device Name**. Select **Next**.

Step 8 Choose an option depending on whether the device is logged into and configured for a manager:

- If your device is brand new and has never been configured for a manager, click **Yes, this new device has never been logged into or configured for a manager**.
- If your device has been previously registered for a manager or is **still** registered to a manager, click **No, this device has been logged into and configured for a manager**.

Step 9 Click **Next**.

Step 10 In the **Policy Assignment** step, use the drop-down menu to select an access control policy to deploy once the device is onboarded. If you have no policies configured, select the **Default Access Control Policy**.

Step 11 Select the subscription licenses you want to apply to the device. Click **Next**.

Step 12 Once the device is synchronized, select the device you just onboarded from the **Security Devices** page and select any of the options listed under the **Device Management** pane located to the right. We strongly recommend these actions:

- If you did not already, create a custom access control policy to customize the security for your environment.

  For more information, see Access Control Overview.
- You can enable Cisco Security Analytics and Logging to view events in the Security Cloud Control dashboard, or register the device to a Secure Firewall Management Center for security analytics.

  For more information, see Cisco Security Analytics and Logging.

Onboard a Firewall Threat Defense Device to On-Prem Firewall Management Center using Zero-Touch Provisioning

Only the Firepower 1000, Firepower 2100, Secure Firewall 1200, and Secure Firewall 3100 devices can be onboarded to on-premises Firewall Management Center using the zero-touch provisioning method.

Before you begin

Confirm the following is completed prior to onboarding:

- You have a Security Cloud Control tenant. If you do not, see Request a Security Cloud Control Tenant for more information.
- Before onboarding any new devices, ensure an on-prem FMC is fully set up, configured, and recognized as a management center within your Security Cloud Control tenant.
- The device is freshly installed and has never been logged into through the device CLI, a Firewall Management Center, or the Firewall Device Manager.

- The device is running version 7.2 or later. .

Procedure

Step 1 In the left pane, click **Security Devices**.

Step 2 Click **Onboard device or service** ( ) icon.

Step 3 Click the **FTD** tile.

Step 4 Under **Management Mode**, select **FTD**. By selecting **FTD** under **Management Mode**, you will not be able to manage the device using the previous management platform. All existing policy configurations except for interface configurations will be reset. You must re-configure policies after you onboard the device.

Step 5 Click **Use Serial Number**.

Step 6 Select an available on-prem FMC from the drop-down list. Click **Next**.

Note

- On-prem FMCs running version 7.4 or later and onboarded with Cisco Security Cloud are displayed in the drop-down.
- Provide the **Public IP address** or **FQDN** value of the selected on-prem FMC unless
  - The FTD is publicly reachable
  - The FTD is running a version earlier than 7.4
  - The connection is being made through the data interface

For information on onboarding a threat defense device to a Cloud-Delivered Firewall Management Center, see Onboard a Threat Defense Device to Cloud-delivered Firewall Management Center using Zero-Touch Provisioning.

Step 7 Enter the **Device Serial Number** and the **Device Name**. Click **Next**.

Step 8 Choose an option depending on whether the device is logged into and configured for a manager:

- If your device is brand new and has never been configured for a manager, click **Yes, this new device has never been logged into or configured for a manager**.
- If your device has been previously registered for a manager or is **still** registered to a manager, click **No, this device has been logged into and configured for a manager**.

Step 9 Click **Next**.

Step 10 In the **Policy Assignment** step, use the drop-down menu to select an access control policy to deploy once the device is onboarded. If you have no policies configured, select the **Default Access Control Policy**.

Step 11 Select the subscription licenses you want to apply to the device. Click **Next**.

Step 12 Once the device is synchronized, select the device you just onboarded from the **Security Devices** page and select any of the options listed under the **Device Management** pane located to the right. We strongly recommend these actions:

- If you did not already, create a custom access control policy to customize the security for your environment.

  For more information, see Access Control Overview.

- You can enable Cisco Security Analytics and Logging to view events in the Security Cloud Control dashboard, or register the device to a Secure Firewall Management Center for security analytics.

  For more information, see Cisco Security Analytics and Logging.

Onboard Firewall Threat Defense Devices using Device Templates to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning

Only the Firepower 1000, Firepower 2100, Secure Firewall 1200, Secure Firewall 200, and Secure Firewall 3100 devices can be onboarded with the zero-touch provisioning method.

Before you begin

Confirm that the following is completed before onboarding:

- You have a Security Cloud Control tenant. If you do not, see Request a Security Cloud Control Tenant for more information.
- Cloud-Delivered Firewall Management Center is enabled for your tenant.
- The device is freshly installed but has never been logged into by either the device CLI, a Firewall Management Center, or the Firewall Device Manager.
- The device is running version 7.4 or later.

Procedure

Step 1 In the left pane, click **Security Devices**.

Step 2 Click **Onboard device or service** ( ) icon.

Step 3 Click the **FTD** tile.

Step 4 Under **Management Mode**, select **FTD**. By selecting **FTD** under **Management Mode**, you will not be able to manage the device using the previous management platform. All existing policy configurations except for interface configurations will be reset. You must re-configure policies after you onboard the device.

Step 5 Click the **Bulk Onboard using CSV File** tile.

Step 6 In the **Template Assignment** field, select a template from the drop-down list. The access control policy associated with the template and the supported device models for the selected template are then displayed. Click **Next**.

Step 7 In the **Upload CSV File** field, **Drag & drop** your CSV template file or **Click** to select the CSV template file that you want to upload.

You can download a **CSV Sample Template File** to have a look at the required header details that have to be used in the template. The CSV template file must be less than 2 MB in size. The filename must satisfy the following criteria:

- Can have a maximum of 64 characters.
- Only alphanumeric characters and special characters such as dash (-), period (.), and underscore (_) are allowed.
- Must not contain any spaces.

A properly formatted .csv file has the following fields:

- Mandatory fields
  - Display Name - Name of the device. Type: string. Example: test1
  - Serial Number - Serial number of the device. Type: string, Example: JADX345670EG
- Optional fields
  - Device Group - Name of the device group, Type: string, Example: testgroup
  - Admin Password - Password for admin access, Type: string, Example: E28@2OiUrhx
- Variables - Use the following format: $<*varName*>. Sample variable: $LAN-Devices-IPv4Address - IPv4 address of the LAN device. Type: string. Example: 1.2.3.4.
- Network object overrides - Use the following format:<*objType*>:<*objName*>. Sample network object override: Network:LAN-Devices-Network - IP address of the network of LAN devices. Type: string. Example: 1.2.3.4.

A sample CSV template file containing configuration for two devices is as given below.

DisplayName	SerialNumber	AdminPassword	$WANLinkIP	Host:gateway
Branch A FTD	JADX345410AB	C15c05n0rt#	10.20.30.1	10.2.3.1
Branch B FTD	JADX345670CE	Admin123!	10.20.30.5	10.2.3.1

Step 8 Click **Next**.

Step 9 A statement mentioning that the devices are being onboarded is displayed. You can check the onboarding status of the devices in the **Security Devices** window.

Once the device is synchronized, select the device you just onboarded from the **Security Devices** page and select any of the options listed under the **Device Management** pane located to the right. We strongly recommend the following actions:

- If you did not already, create a custom access control policy to customize the security for your environment. See Access Control Overview in *Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control* for more information.
- Enable Cisco Security Analytics and Logging (SAL) to view events in the Security Cloud Control dashboard **or** register the device to an Secure Firewall Management Center for security analytics. See

Cisco Security Analytics and Logging in *Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control* for more information.

Cloud Deployment

Cloud Deployment helps you deploy cloud-based enforcement points based on your intended outcomes and use cases. This provides the most cost-effective, scalable, and fully featured security for your multicloud needs.

With Cloud Deployment, you can:

- deploy faster with a guided, intuitive interface,
- select and launch the right enforcement point for your use case, and
- scale automatically to match demand and sustain performance.

Available enforcement points:

- Cisco Secure Firewall Threat Defense Virtual (FTDv): recommended for TLS Decrypt, IDS/IPS, Antivirus, Geo IP, or Malicious IP.
- Multicloud Defense: recommended for Proxy, WAF – L7 DOS, DLP, or Two-Core Enforcement.

Supported cloud service providers are:

- AWS
- Azure
- Google Cloud Platform (GCP)
- Oracle Cloud Infrastructure (OCI): Supported in this release by Multicloud Defense gateway only

Configure FTDv enforcement point

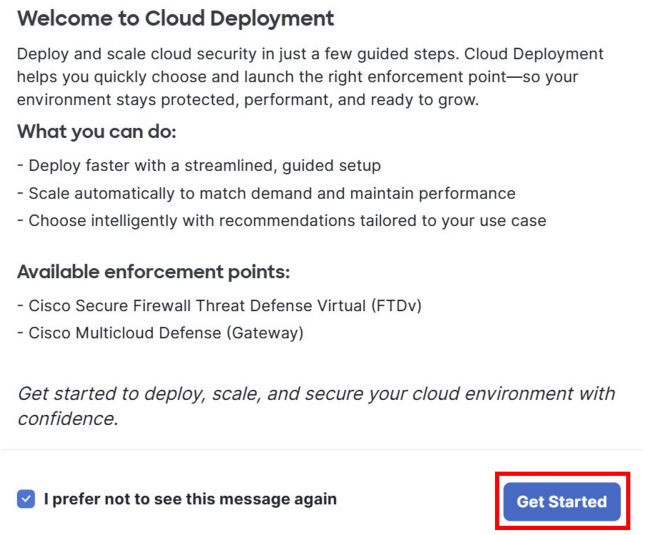
Cloud Deployment helps you deploy cloud-based enforcement points based on your intended outcomes and use cases. If you select TLS Decrypt, IDS/IPS, Antivirus, Geo IP, or Malicious IP, Cloud Deployment uses the Cisco Secure Firewall Threat Defense Virtual (FTDv) gateway for deployment.

Follow these steps to configure your FTDv enforcement point:

Procedure

Step 1 Choose **Security Devices** > **Cloud Deployment**.

Step 2 Click **Get Started**.



Step 3 On the **Account & Enforcement Type** tab, enter these details.

Field	Description
Account	Select an account from the drop-down list: AWS, Azure, or Google Cloud Platform (GCP).
Region	Select the region this gateway will be deployed into.
Select the type of Traffic Flow	Select the type of traffic flow (at least one must be selected): Ingress: Traffic entering from external destination Egress: Traffic leaving a network to external destination East-West: Traffic communicated across workloads with the network
Select your architectural requirements	Select TLS Decrypt, IDS/IPS, Antivirus, Geo IP, or Malicious IP to deploy through the FTDv gateway. Select Proxy, WAF – L7 DOS, DLP, or Two-Core Enforcement to deploy through Multicloud Defense gateway. Note that if you select Two-Core Enforcement, Antivirus is unavailable as it is not supported.

Step 4 Click **Next**.

Step 5 On the **Configure Enforcement Point** tab, enter these details.

Firepower Threat Defense Virtual Enforcement Information	
Name	Enter a unique name for this enforcement point; the name can be up to 53 characters long and must be alphanumeric and include an underscore or a hyphen.

Description	(Optional) Enter a description for this enforcement point.

If you selected AWS account earlier, enter these details:

VPC/VNet	You can select an existing Service VPC/VNet from the drop-down list or click **+ Create SVPC** to create a new one. After creation and activation, the wizard auto-populates the related network fields. This step can take up to three to four minutes.
Mgmt. Security Group	Select the security group to associate with the management interface.
Datapath Security Group 1	Select the security group to associate with datapath 1 interface.
Datapath Security Group 2	Select the security group to associate with datapath 2 interface.
Availability Zone	(Optional) Select the specific availability zone within a cloud region where a particular FTDv gateway instance is deployed. Click **+ Add Availability Zone** to add a new availability zone.
Management Subnet	(Optional) Select the subnet for the management interface in the selected availability zone.
Datapath Subnet 1	(Optional) Select the subnet for the primary datapath interface in the selected availability zone.
Datapath Subnet 2	(Optional) Select the subnet for the secondary datapath interface.
Cisco Secure Firewall Threat Defense Virtual	
License Model	Select the licensing model for the deployment: **Smart Licensing**, **Gateway Hours**, or **Pay As You Go**
Smart Licensing/Gateway Hours/Pay As You Go	
Software Version	Select the FTDv release to deploy.
Admin Password	Enter the Admin password. It must be between 12 and 72 characters in length and include uppercase and lowercase letters, numeric digits, and special characters, excluding the backslash (\) and hyphen (-) characters. The password cannot contain more than two identical or sequential characters in a row (such as 'aaa' or '123').
Policy Ruleset	Select the FMC policy to attach to the FTDv deployment.
Performance Tier	(Smart Licensing only) Expand the drop-down menu and select the appropriate performance tier for your device. FTDv50 is selected by default.

License Types	(Smart Licensing only) Expand the drop-down menu and select the appropriate license type that you have purchased or will purchase in the future. The **Base** license is auto-selected. For more information about different licensing types, refer to Licensing. If you want to enable malware defense features, you must also have the THREAT license active. The THREAT license is required to support the MALWARE license functionalities.
Gateway Configuration	
Instance Type	Select the VM size used for the gateway instances.
Minimum Instances	(Optional) Enter the minimum number of instances that you plan to deploy. This must be at least 1.
Maximum Instances	(Optional) Enter the maximum number instances that you plan to deploy. This is the maximum number that is used for auto-scaling in each availability zone. This number must be greater than or equal to minimum instances.
Health Check Port	(Optional) Enter the port number that the load balancer uses to check instance health. Datapath security groups assigned to the instance(s) must allow traffic on this port. The default port is 65534, and the valid range is from 1 to 65535.
Cloud Configuration	
Key Pair	Select the key pair that is associated with the cloud account you selected earlier. You must select an existing key pair from the chosen region.
EBS Encryption	Select the appropriate EBS encryption for your AWS account.
Advanced Settings	
Disable Public IP	Check the **Disable Public IP** check box. Disabling this setting reduces the exposure of your network to external threats. Private IP addresses help protect against unauthorized access, potential attacks, and provide better control over internal traffic. This setting applies only when using NAT-based service and only for data interfaces.

If you selected Azure account earlier, enter these details:

VPC/VNet	You can select an existing Service VPC/VNet from the drop-down list or click **+ Create SVPC** to create a new one. After creation and activation, the wizard auto-populates the related network fields. This step can take up to three to four minutes.

Resource Group	Select the resource group that owns the deployment resources.
Mgmt. Security Group	Select the security group to associate with the management interface.
Datapath Security Group 1	Select the security group to associate with datapath 1 interface.
Datapath Security Group 2	Select the security group to associate with datapath 2 interface.
Availability Zone	(Optional) Select the specific availability zone within a cloud region where a particular FTDv gateway instance is deployed. Click **+ Add Availability Zone** to add a new availability zone.
Management Subnet	(Optional) Select the subnet for the management interface in the selected availability zone.
Datapath Subnet 1	(Optional) Select the subnet for the primary datapath interface in the selected availability zone.
Datapath Subnet 2	(Optional) Select the subnet for the secondary datapath interface.
Cisco Secure Firewall Threat Defense Virtual	
License Model	Select the licensing model for the deployment: **Smart Licensing** or **Gateway Hours**
Smart Licensing/Gateway Hours	
Software Version	Select the FTDv release to deploy.
Admin Password	Enter the Admin password. It must be between 12 and 72 characters in length and include uppercase and lowercase letters, numeric digits, and special characters, excluding the backslash (\) and hyphen (-) characters. The password cannot contain more than two identical or sequential characters in a row (such as 'aaa' or '123').
Policy Ruleset	Select the FMC policy to attach to the FTDv deployment.
Performance Tier	(Smart Licensing only) Expand the drop-down menu and select the appropriate performance tier for your device. FTDv50 is selected by default.

License Types	(Smart Licensing only) Expand the drop-down menu and select the appropriate license type that you have purchased or will purchase in the future. The **Base** license is auto-selected. For more information about different licensing types, refer to Licensing. If you want to enable malware defense features, you must also have the THREAT license active. The THREAT license is required to support the MALWARE license functionalities.
Gateway Configuration	
Instance Type	Select the VM size used for the gateway instances.
Minimum Instances	(Optional) Enter the minimum number of instances that you plan to deploy. This must be at least 1.
Maximum Instances	(Optional) Enter the maximum number instances that you plan to deploy. This is the maximum number that is used for auto-scaling in each availability zone. This number must be greater than or equal to minimum instances.
Health Check Port	(Optional) Enter the port number that the load balancer uses to check instance health. Datapath security groups assigned to the instance(s) must allow traffic on this port. The default port is 65534, and the valid range is from 1 to 65535.
Azure Configuration	
Key Selection	You can either provide an **SSH Public Key** directly or select an existing **SSH Key Pair** from the drop-down list.
Disk Encryption	Select either a provider-managed encryption key or a customer-managed key (CMK). If using a CMK, provide the **Disk Encryption Set ID**.
Advanced Settings	
Disable Public IP	Check the **Disable Public IP** check box. Disabling this setting reduces the exposure of your network to external threats. Private IP addresses help protect against unauthorized access, potential attacks, and provide better control over internal traffic. This setting applies only when using NAT-based service and only for data interfaces.

If you selected GCP account earlier, enter these details:

GCP Configuration

Datapath VPC	You can select an existing Service VPC/VNet from the drop-down list or click **+ Create SVPC** to create a new one. After creation and activation, the wizard auto-populates the related network fields. This step can take up to three to four minutes.
Datapath Network Tag 1	Select the network tag used for datapath firewall rules.
Datapath Network Tag 2	This is the second datapath tag used for dual-interface FTDv layouts. This identifies a separate management VPC.
Management VPC	Select the management VPC from the drop-down list.
Management Network Tag	Select the network tag used for management firewall rules from the drop-down list.
SSH Public Key	(Optional) Paste the SSH public key.
Disk Encryption	Select either a provider-managed encryption key or a customer-managed key (CMK). If using a CMK, provide the **Cloud KMS key resource name**.
Availability Zone	(Optional) Select the specific availability zone within a cloud region where a particular FTDv gateway instance is deployed. Click **+ Add Availability Zone** to add a new availability zone.
Management Subnet	(Optional) Select the subnet for the management interface in the selected availability zone.
Datapath Subnet 1	(Optional) Select the subnet for the primary datapath interface in the selected availability zone.
Datapath Subnet 2	(Optional) Select the subnet for the secondary datapath interface.
Cisco Secure Firewall Threat Defense Virtual	
License Model	Select the licensing model for the deployment: **Smart Licensing** or **Gateway Hours**
Smart Licensing/Gateway Hours	
Software Version	Select the FTDv release to deploy.
Admin Password	Enter the Admin password. It must be between 12 and 72 characters in length and include uppercase and lowercase letters, numeric digits, and special characters, excluding the backslash (\) and hyphen (-) characters. The password cannot contain more than two identical or sequential characters in a row (such as 'aaa' or '123').
Policy Ruleset	Select the FMC policy to attach to the FTDv deployment.

Performance Tier	(Smart Licensing only) Expand the drop-down menu and select the appropriate performance tier for your device. FTDv50 is selected by default.
License Types	(Smart Licensing only) Expand the drop-down menu and select the appropriate license type that you have purchased or will purchase in the future. The **Base** license is auto-selected. For more information about different licensing types, refer to Licensing. If you want to enable malware defense features, you must also have the THREAT license active. The THREAT license is required to support the MALWARE license functionalities.
Gateway Configuration	
Instance Type	Select the VM size used for the gateway instances.
Minimum Instances	(Optional) Enter the minimum number of instances that you plan to deploy. This must be at least 1.
Maximum Instances	(Optional) Enter the maximum number instances that you plan to deploy. This is the maximum number that is used for auto-scaling in each availability zone. This number must be greater than or equal to minimum instances.
Health Check Port	(Optional) Enter the port number that the load balancer uses to check instance health. Datapath security groups assigned to the instance(s) must allow traffic on this port. The default port is 65534, and the valid range is from 1 to 65535.
Advanced Settings	
Disable Public IP	Check the **Disable Public IP** check box. Disabling this setting reduces the exposure of your network to external threats. Private IP addresses help protect against unauthorized access, potential attacks, and provide better control over internal traffic. This setting applies only when using NAT-based service and only for data interfaces.

Step 6 Click **Next**. If you want to modify the settings, click **Back**

Step 7 On the **Review & Deploy** tab, review your configuration and deploy the enforcement point. If you encounter an error, return to the **Configure Enforcement Points** tab to correct it.

Step 8 Click **Deploy**.

Multicloud Defense deploys the gateway.

You must attach at least one ruleset to the gateway before you secure a spoke VPC/VNet. For more information about rule sets, refer to Rule Sets and Rule Set Groups.

Configure Multicloud Defense enforcement point

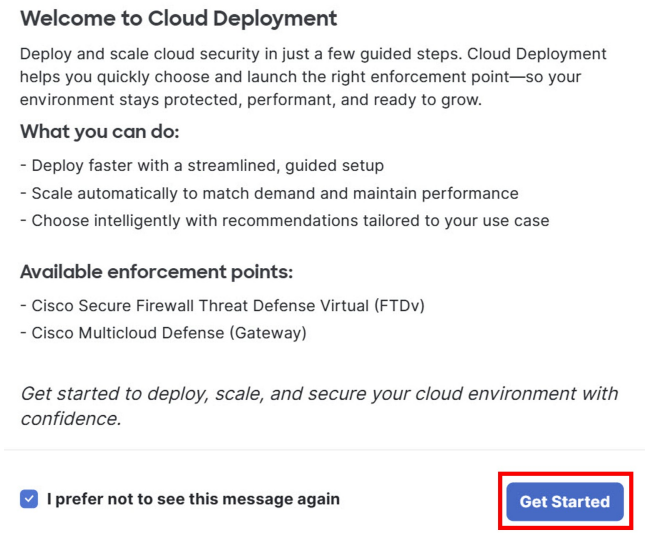
Cloud Deployment helps you deploy cloud-based enforcement points based on your intended outcomes and use cases. If you select Proxy, WAF – L7 DOS, DLP, or Two-Core Enforcement, Cloud Deployment uses the Multicloud Defense gateway for deployment.

Follow these steps to configure Multicloud Defense enforcement point:

Procedure

Step 1 Choose **Security Devices** > **Cloud Deployment**.

Step 2 Click **Get Started**.



Step 3 On the **Account & Enforcement Type** tab, enter these details.

Field	Description
Account	Select an account from the drop-down list: AWS, Azure, Google Cloud Platform (GCP), or Oracle Cloud Infrastructure (OCI).
Region	Select the region this gateway will be deployed into.
Select the type of Traffic Flow	Select the type of traffic flow (at least one must be selected): Ingress: Traffic entering from an external destination Egress: Traffic leaving a network to external destination East-West: Traffic communicated across workloads with the network

Select your architectural requirements	Select TLS Decrypt, IDS/IPS, Antivirus, Geo IP, or Malicious IP if you want to deploy these features through the FTDv gateway. Select Proxy, WAF – L7 DOS, DLP, or Two-Core Enforcement if you want to deploy through the Multicloud Defense gateway. Note that if you select Two-Core Enforcement, Antivirus is unavailable as it is not supported.

Step 4 Click **Next**.

Step 5 On the **Configure Enforcement Point** tab, enter these details.

Multicloud Defense Gateway Enforcement Information	
Name	Enter a unique name for this enforcement point; the name can be up to 53 characters long and must be alphanumeric and include an underscore or a hyphen.
Description	(Optional) Enter a description for this enforcement point.

If you selected AWS account earlier, enter these details:

Multicloud Defense Gateway Enforcement Information	
Deployment	Click **Centralized** or **Distributed** to select the deployment type. **Centralized** allows you to select only a Service VPC. **Distributed** allows you to select only a regular VPC/VNet. If you select a Service VPC, the deployment operates in Hub mode. If you select a regular VPC or VNet, the deployment operates in Edge mode.
VPC/VNet	You can select an existing Service VPC/VNet from the drop-down list or click **+ Create SVPC** to create a new one. You cannot create a new Service VPC/VNet if you selected **Distributed** earlier. After you create and activate the service, the wizard automatically fills in the related network fields. This step can take up to 3 to 4 minutes.
Mgmt. Security Group	Select the security group to associate with the management interface.
Datapath Security Group	Select the security group to associate with the datapath interface.

Availability Zone	(Optional) Select the specific availability zone within a cloud region where a particular Multicloud Defense gateway instance is deployed. Click **+ Add Availability Zone** to add a new availability zone.
Management Subnet	(Optional) Select the subnet for the management interface in the selected availability zone.
Datapath Subnet	(Optional) Select the subnet for the primary datapath interface in the selected availability zone.
Gateway Configuration	
Instance Type	Select the VM size used for the gateway instances from the drop-down list.
Minimum Instances	(Optional) Enter the minimum number of instances you plan to deploy. This must be at least one.
Maximum Instances	(Optional) Enter the maximum number of instances you plan to deploy. This is the maximum number that is used for autoscaling in each availability zone. This number must be greater than or equal to minimum instances.
Health Check Port	Enter the port number used by Multicloud Defense load balancer to check the health of the instances. Datapath security groups assigned to the instance(s) must allow traffic on this port. The default port is 65534, and the valid range is from 1 to 65535.
Gateway Image Group	Select the logical grouping of available gateway images from the drop-down list.
Gateway Image	Select the exact gateway image/version to deploy from the drop-down list.
Policy Ruleset	Select the Multicloud Defense policy ruleset to attach to the gateway.
Cloud Configuration	
Key Pair	Select the key pair associated with the cloud account you selected in the previous screen.
Gateway IAM Role	Select the IAM role that allows the gateway to perform READ and WRITE operations on your AWS account. Multicloud Defense creates this role for you when you save and deploy the gateway.
EBS Encryption	Select the appropriate EBS encryption for your specific AWS account.
Multicloud Defense Profiles	
Packet Capture Profile	(Optional) Select the packet capture profile for threat and flow PCAPs.

Log Forwarding Profile	(Optional) Select the log forwarding profile used to forward Events/Logs to a security information and event management (SIEM) system.
Metrics Profile	(Optional) Select the metrics profile.
NTP Profile	(Optional) Select the Network Time Protocol (NTP) for time synchronization.
Advanced Settings	
Use Global Accelerator	(Ingress only) Associate the gateway with an AWS Global Accelerator listener so that traffic uses your accelerator.
Disable Public IP	By default, the Multicloud Defense gateway enables the use of the public IP of the router available. If you do not want this enabled, check the **Disable Public IP** check box. This reduces the exposure of your network to external threats.
Use Internal Load Balancer	To reduce the exposure of your network to external threats, you can use an internal load balancer instead of a public one for the gateway front end (Ingress only). While the Multicloud Defense gateway enables a public router IP by default, you can opt for a private configuration by checking the **Disable Public IP** check box.

If you selected Azure account earlier, enter these details:

Multicloud Defense Gateway Enforcement Information	
Deployment	Click **Centralized** or **Distributed** to select the deployment type. **Centralized** allows the selection of only a Service VPC. **Distributed** allows the selection of only a regular VPC/VNet. If you select a Service VPC, the deployment operates in Hub mode. If you select a regular VPC or VNet, the deployment operates in Edge mode.
VPC/VNet	You can select an existing Service VPC/VNet from the drop-down list or click **+ Create SVPC** to create a new one. You cannot create a new Service VPC/VNet if you selected **Distributed** earlier. After you create and activate the service, the wizard automatically fills in the related network fields. This step can take up to 3 to 4 minutes.
Resource Group	Select the resource group that owns the deployment resources.

Mgmt. Security Group	Select the security group to associate with the management interface.
Datapath Security Group 1	Select the security group to associate with the datapath interface.
Availability Zone	(Optional) Select the specific availability zone within a cloud region where a particular Multicloud Defense gateway instance is deployed. Click **+ Add Availability Zone** to add a new availability zone.
Management Subnet	(Optional) Select the subnet for the management interface in the selected availability zone.
Datapath Subnet	(Optional) Select the subnet for the primary datapath interface in the selected availability zone.
Gateway Configuration	
Instance Type	Select the VM size used for the gateway instances.
Minimum Instances	(Optional) Enter the minimum number of instances you plan to deploy. This must be at least one.
Maximum Instances	(Optional) Enter the maximum number of instances you plan to deploy. This is the maximum number that is used for autoscaling in each availability zone. This number must be greater than or equal to minimum instances.
Health Check Port	Enter the port number that the load balancer uses to check instance health. Datapath security groups assigned to the instance(s) must allow traffic on this port. The default port is 65534, and the valid range is from 1 to 65535.
Gateway Image Group	Select the logical grouping of available gateway images from the drop-down list.
Gateway Image	Select the exact gateway image/version to deploy from the drop-down list.
Policy Ruleset	Select the Multicloud Defense policy ruleset to attach to the gateway.
Azure Configuration	
Key Selection	You can either provide an **SSH Public Key** directly or select an existing **SSH Key Pair** from the drop-down list.
Username	The default SSH username is centos.
User Assigned Identity ID	(Optional) Used for accessing Key Vault or writing PCAP files to Blob Storage.
Disk Encryption	(Optional) Select the disk encryption type from the drop-down list

Multicloud Defense Profiles	
Packet Capture Profile	(Optional) Select the packet capture profile for threat and flow PCAPs.
Log Forwarding Profile	(Optional) Select the log forwarding profile used to forward Events/Logs to a security information and event management (SIEM) system.
Metrics Profile	(Optional) Select the metrics profile.
NTP Profile	(Optional) Select the Network Time Protocol (NTP) for time synchronization.
Advanced Settings	
Management DNS Server	(Optional) Enter the IP address of a DNS server reachable by the gateway for management traffic.
Attach Additional Load Balancer	(Optional) Click **+ Add** to associate Azure Load Balancers with this gateway for advanced traffic distribution. Select the **Azure LB Name** from the drop-down list and enter the **Backend Pool Name** and **Resource Group**.
Disable Public IP	By default, the Multicloud Defense gateway enables the use of the public IP of the router available. If you do not want this enabled, check the **Disable Public IP** check box. This reduces the exposure of your network to external threats.
Use Internal Load Balancer	To reduce the exposure of your network to external threats, you can use an internal load balancer instead of a public one for the gateway front end (Ingress only). While the Multicloud Defense gateway enables a public router IP by default, you can opt for a private configuration by checking the **Disable Public IP** check box.

If you selected GCP account earlier, enter these details:

Deployment	Click **Centralized** or **Distributed** to select the deployment type. **Centralized** allows the selection of only a Service VPC. **Distributed** allows the selection of only a regular VPC/VNet. If you select a Service VPC, the deployment operates in Hub mode. If you select a regular VPC or VNet, the deployment operates in Edge mode.
GCP Configuration	
Service Account Email	Enter the service account email attached to gateway instances.

Datapath VPC	You can select an existing Service VPC/VNet from the drop-down list or click **+ Create SVPC** to create a new one. After you create and activate the service, the wizard automatically fills in the related network fields. This step can take up to 3 to 4 minutes.
Datapath Network Tag	Select the network tag used for datapath firewall rules.
Management VPC	Select the management VPC from the drop-down list.
Management Network Tag	Select the network tag used for management firewall rules from the drop-down list.
Username	The default SSH username is centos.
SSH Public Key	(Optional) Paste the SSH public key.
Disk Encryption	(Optional) Select the disk encryption type from the drop-down list
Network Details	
Availability Zone	(Optional) Select the specific availability zone within a cloud region where a particular Multicloud Defense gateway instance is deployed. Click **+ Add Availability Zone** to add a new availability zone.
Management Subnet	(Optional) Select the subnet for the management interface in the selected availability zone.
Datapath Subnet	(Optional) Select the subnet for datapath interface in the selected availability zone.
Gateway Configuration	
Instance Type	Select the VM size used for the gateway instances.
Minimum Instances	(Optional) Enter the minimum number of instances you plan to deploy. This must be at least one.
Maximum Instances	(Optional) Enter the maximum number of instances you plan to deploy. This is the maximum number that is used for autoscaling in each availability zone. This number must be greater than or equal to minimum instances.
Health Check Port	Enter the port number that the load balancer uses to check instance health. Datapath security groups assigned to the instance(s) must allow traffic on this port. The default port is 65534, and the valid range is from 1 to 65535.
Gateway Image Group	Select the logical grouping of available gateway images from the drop-down list.
Gateway Image	Select the exact gateway image/version to deploy from the drop-down list.

Policy Ruleset	Select the Multicloud Defense policy ruleset to attach to the gateway.
Multicloud Defense Profiles	
Packet Capture Profile	(Optional) Select the packet capture profile for threat and flow PCAPs.
Log Forwarding Profile	(Optional) Select the log forwarding profile used to forward Events/Logs to a security information and event management (SIEM) system.
Metrics Profile	(Optional) Select the metrics profile.
NTP Profile	(Optional) Select the Network Time Protocol (NTP) for time synchronization.
Advanced Settings	
Disable Public IP	By default, the Multicloud Defense gateway enables the use of the public IP of the router available. If you do not want this enabled, check the **Disable Public IP** check box. This reduces the exposure of your network to external threats.
GCP LB Protocol UDP	Check this check box to enable UDP on the gateway's GCP load balancer configuration.

If you selected Oracle Cloud Infrastructure (OCI) account earlier, enter these details:

Multicloud Defense Gateway Enforcement Information	
Deployment	Click **Centralized** or **Distributed** to select the deployment type. **Centralized** allows the selection of only a Service VPC. **Distributed** allows the selection of only a regular VPC/VNet. If you select a Service VPC, the deployment operates in Hub mode. If you select a regular VPC or VNet, the deployment operates in Edge mode.
VPC/VNet	You can select an existing Service VPC/VNet from the drop-down list or click **+ Create SVPC** to create a new one. You cannot create a new Service VPC/VNet if you selected **Distributed** earlier. After you create and activate the service, the wizard automatically fills in the related network fields. This step can take up to 3 to 4 minutes.
Mgmt. Security Group	(Optional) Select the security group to associate with the management interface.

Datapath Security Group	(Optional) Select the security group to associate with the datapath interface.
Availability Zone	(Optional) Select the specific availability zone within a cloud region where a particular Multicloud Defense gateway instance is deployed. Click **+ Add Availability Zone** to add a new availability zone.
Management Subnet	(Optional) Select the subnet for the management interface in the selected availability zone.
Datapath Subnet	(Optional) Select the subnet for the primary datapath interface in the selected availability zone.
Gateway Configuration	
Instance Type	Select the VM size used for the gateway instances from the drop-down list.
Minimum Instances	(Optional) Enter the minimum number of instances you plan to deploy. This must be at least one.
Maximum Instances	(Optional) Enter the maximum number of instances you plan to deploy. This is the maximum number that is used for autoscaling in each availability zone. This number must be greater than or equal to minimum instances.
Health Check Port	Enter the port number used by Multicloud Defense load balancer to check the health of the instances. Datapath security groups assigned to the instance(s) must allow traffic on this port. The default port is 65534, and the valid range is from 1 to 65535.
Gateway Image Group	Select the logical grouping of available gateway images from the drop-down list.
Gateway Image	Select the exact gateway image/version to deploy from the drop-down list.
Policy Ruleset	Select the Multicloud Defense policy ruleset to attach to the gateway.
OCI Configuration	
Compartment	OCI Compartment where Multicloud Defense gateway will be deployed.
SSH Public Key	Paste the SSH public key.
Multicloud Defense Profiles	
Packet Capture Profile	Select the packet capture profile for threat and flow PCAPs.
Log Forwarding Profile	Select the log forwarding profile used to forward Events/Logs to a security information and event management (SIEM) system.
Metrics Profile	Select the metrics profile.

NTP Profile	Select the Network Time Protocol (NTP) for time synchronization.
Advanced Settings	
Internal Load Balancer	To reduce the exposure of your network to external threats, you can use an internal load balancer instead of a public one for the gateway front end (Ingress only). While the Multicloud Defense gateway enables a public router IP by default, you can opt for a private configuration by selecting the **Disable Public IP** check box.
User ID	Enable User ID on the gateway if your account supports it.
Management DNS Server	Enter the IP address of a DNS server reachable by the gateway for management traffic.
Public IP	By default, the Multicloud Defense gateway enables the use of the public IP of the router available. If you do not want this enabled, check the **Disable Public IP** check box. This reduces the exposure of your network to external threats.

Step 6 Click **Next**. If you want to modify the settings, click **Back**.

Step 7 On the **Review & Deploy** tab, review your configuration and deploy the enforcement point. If you encounter an error, return to the **Configure Enforcement Points** tab to correct it.

Step 8 Click **Deploy**.

Multicloud Defense deploys the gateway.

You must attach at least one ruleset to the gateway before you secure a spoke VPC/VNet. For more information about rule sets, refer to Rule Sets and Rule Set Groups.

Onboard a Secure Firewall Threat Defense Cluster

This task enables you to register a Secure Firewall Threat Defense cluster with Cloud-Delivered Firewall Management Center for centralized management and policy enforcement.

The following table provides information about device models that support cluster onboarding and creation on the Cloud-Delivered Firewall Management Center:

Secure Firewall Threat Defense Platforms	Minimum Secure Firewall Threat Defense Version for Cluster Management	Support cluster creation from Cloud-Delivered Firewall Management Center?
VMware, KVM	7.2.1	Yes
AWS, GCP	7.2.1	No
Azure	7.3	No
Secure Firewall 3100	7.2.1	Yes

Secure Firewall Threat Defense Platforms	Minimum Secure Firewall Threat Defense Version for Cluster Management	Support cluster creation from Cloud-Delivered Firewall Management Center?
Firepower 4100	7.2.x	No
Secure Firewall 4200	7.4	Yes
Firepower 9300	7.2.x	No



Note When you onboard a Secure Firewall Threat Defense cluster to Cloud-Delivered Firewall Management Center, you register only the cluster CONTROL NODE. After the CONTROL NODE registers, Cloud-Delivered Firewall Management Center discovers the DATA nodes in the cluster.

Before you begin

Review these limitations:

- Firepower 4100 and Firepower 9300 devices must be clustered through the device's Firewall Chassis Manager.
- Secure Firewall 3100 devices, Secure Firewall 4200 devices, KVM, and VMware environments must be clustered through the Cloud-Delivered Firewall Management Center UI.
- Azure, AWS, and GCP environment clusters must be created through their own environment and onboarded to Cloud-Delivered Firewall Management Center.
- Verify that the Secure Firewall Threat Defense virtual devices are successfully configured as a cluster before you register the cluster with Cloud-Delivered Firewall Management Center.

  For example, in the Secure Firewall Threat Defense CLI, enter the `show cluster info` command. Confirm that the output shows the cluster as `On`, one unit in `CONTROL_NODE` state, and the remaining units in `DATA_NODE` state.

```
> show cluster info
Cluster mAZ-ngfw-cl: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "74-a" in state DATA_NODE
        ID        : 2
        Version   : 9.22(1)1
        Serial No.: 9AUVQ3DSF66
        CCL IP    : 1.1.1.74
        CCL MAC   : 02e2.778f.d3ed
        Module    : NGFWv
        Resource  : 4 cores / 7680 MB RAM
        Last join : 07:28:26 UTC Sep 25 2024
        Last leave: 07:28:11 UTC Sep 25 2024
Other members in the cluster:
    Unit "135-b" in state CONTROL_NODE
        ID        : 0
        Version   : 9.22(1)1
        Serial No.: 9A6W0A51KGK
        CCL IP    : 1.1.2.135
        CCL MAC   : 1294.34ae.4ce9
        Module    : NGFWv
        Resource  : 4 cores / 7680 MB RAM
        Last join : 09:45:52 UTC Sep 24 2024
        Last leave: N/A
    Unit "183-c" in state DATA_NODE
        ID        : 1
        Version   : 9.22(1)1
        Serial No.: 9A1S400HL8F
        CCL IP    : 1.1.3.183
        CCL MAC   : 0aff.e889.f193
        Module    : NGFWv
        Resource  : 4 cores / 7680 MB RAM
        Last join : 07:29:29 UTC Sep 25 2024
        Last leave: 07:28:11 UTC Sep 25 2024
>
```

Follow these steps to onboard a Secure Firewall Threat Defense cluster:

Procedure

Step 1 In the left pane, click **Security Devices**.

Step 2 Click **Onboard device or service** ( ) icon.

Step 3 Click the **FTD** tile.

Step 4 Under **Management Mode**, select **FTD**. By selecting **FTD** under **Management Mode**, you will not be able to manage the device using the previous management platform. All existing policy configurations except for interface configurations will be reset. You must re-configure policies after you onboard the device.

Step 5 Select **Use CLI Registration Key**.

Step 6 Enter the device name in the **Device Name** field and click **Next**.

Step 7 In the **Policy Assignment** step, use the drop-down menu to select an access control policy to deploy once the device is onboarded. If you have no policies configured, select the **Default Access Control Policy**.

Step 8 Specify whether the device you are onboarding is a physical or virtual device. If you are onboarding a virtual device, you must select the device's performance tier from the drop-down menu.

Step 9 Select the subscription licenses you want to apply to the device. Click **Next**.

Step 10 Security Cloud Control generates a command with the registration key. Paste the entire registration key into the device's CLI without making changes.

Step 11 (Optional) Add labels to your device to help sort and filter the **Security Devices** page. Enter a label and select the blue plus button. Labels are applied to the device after it's onboarded to Security Cloud Control.

Once the device is synchronized, select the device you just onboarded from the **Security Devices** page and select any of the options listed under the **Device Management** pane located to the right. We strongly recommend the following actions:

- If you did not already, create a custom access control policy to customize the security for your environment. See Access Control Overview in *Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control* for more information.
- Enable Cisco Security Analytics and Logging (SAL) to view events in the Security Cloud Control dashboard **or** register the device to an Secure Firewall Management Center for security analytics. See Cisco Security Analytics and Logging in *Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control* for more information.

Onboard a Chassis

You can add a Firepower 4100/9300 chassis to the Cloud-Delivered Firewall Management Center. The management center and the chassis share a separate management connection using the chassis MGMT interface. The Cloud-Delivered Firewall Management Center offers chassis-level health alerts. For configuration, you still need to use the Firewall Chassis Manager or FXOS CLI.

Security Cloud Control does not record the processes on the **Workflow** page when you onboard a chassis or provision a multi-instance Firewall Management CenterFirewall Threat Defense device from the Cloud-Delivered Firewall Management Center chassis manager page.

Note For the Secure Firewall 3100/4200, you need to first convert to multi-instance mode. See Onboard the multi-instance chassis using the CLI, on page 667.

Procedure

Step 1 Connect to the chassis FXOS CLI, either using the console port or SSH.

Step 2 In the left pane, click **Security Devices**.

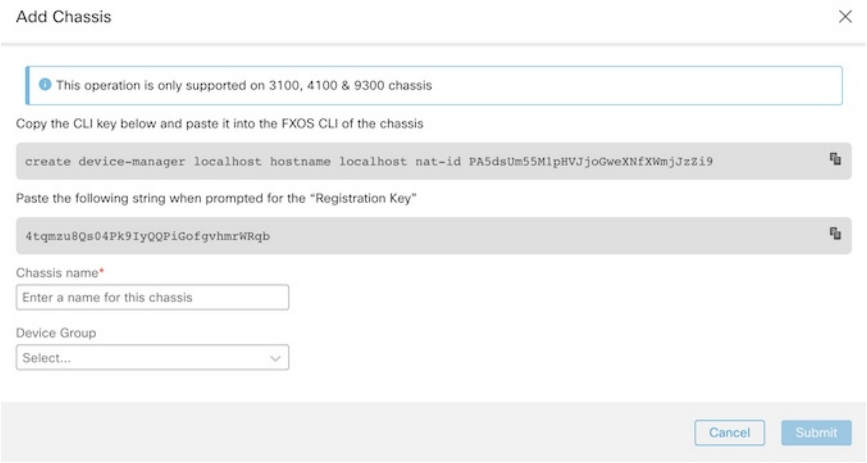
Step 3 Click **Onboard device or service** ( ) icon.

Step 4 Click the **FTD Chassis** tile to open the **Add Chassis** dialog box.

Figure 5: FTD Chassis Tile



Figure 6: Add Chassis



Step 5 Click **Copy** ( ) to copy the top generated command, then paste it at the FXOS CLI of your chassis.

Step 6 When prompted for the Registration Key at the FXOS CLI, click **Copy** ( ) on the **Add Chassis** dialog box for the generated registration key and paste it at the FXOS CLI.

You can disconnect from the FXOS CLI at this point.

Step 7 In the Cloud-Delivered Firewall Management Center **Chassis Name** field, enter a name for the chassis as you want it to display in the Cloud-Delivered Firewall Management Center.

Step 8 (Optional) Add the chassis to a **Device Group**.

Step 9 Click **Submit**.

The chassis is added to the **Device** > **Device Management** page.

Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI

Connect to the device's CLI to perform initial setup, including setting the management IP address, gateway, and other basic networking settings using the setup wizard. Ensure all DNS and firewall ports are accessible for communication.

The dedicated management interface is a special interface with its own network settings. If you do not want to use the management interface, you can use the CLI to configure a data interface instead.

This configuration is ideal for devices that are going to be onboarded with their CLI registration key.

Note Do **not** use this configuration procedure for devices that are onboarding with zero-touch provisioning.

Procedure

Step 1 Connect to the device's CLI, either from the console port or using SSH to the management interface. If you intend to change the network settings, we recommend using the console port so you do not get disconnected.

(Firepower and Secure Firewall hardware models) The console port connects to the FXOS CLI. The SSH session connects directly to the Firewall Threat Defense CLI.

Step 2 Log in with the username **admin** and the password **Admin123** .

At the console port, you connect to the FXOS CLI. The first time you log in to FXOS, you are prompted to change the password. This password is also used for the Firewall Threat Defense login for SSH.

Note
If the password was already changed, and you do not know it, you must reimage the device to reset the password to the default.

For Firepower and Secure Firewall hardware, see the Reimage Procedures in the Cisco FXOS Troubleshooting Guide for the Firewall Threat Defense .

For the ISA 3000, see the Cisco Secure Firewall ASA and Threat Defense Reimage Guide .

Example:

```
firepower login:
                        admin
Password:
                        Admin123
Successful login attempts for user 'admin' : 1
[...]
Hello admin. You must change your password.
Enter new password:
                        ********
Confirm new password:
                        ********
Your password was updated successfully.
```

```
[...]
firepower#
```

Step 3 If you connected to FXOS on the console port, connect to the Firewall Threat Defense CLI.

connect ftd

Example:

```
firepower# connect ftd
>
```

Step 4 The first time you log in to the device, you are prompted to accept the End User License Agreement (EULA) and, if using an SSH connection, to change the admin password. You are then presented with the CLI setup script.

Note
You cannot repeat the CLI setup wizard unless you clear the configuration; for example, by reimaging. However, all of these settings can be changed later at the CLI using **configure network** commands. See the threat defense command reference.

Defaults or previously entered values appear in brackets. To accept previously entered values, press **Enter**.

Note
The management interface settings are used even when you enable Firewall Threat Defense access on a data interface. For example, the management traffic that is routed over the backplane through the data interface will resolve FQDNs using the management interface DNS servers, and not the data interface DNS servers.

See the following guidelines:

- **Configure IPv4 via DHCP or manually?**—If you want to use a data interface for Firewall Threat Defense access instead of the management interface, choose **manual**. Although you do not plan to use the management interface, you must set an IP address, for example, a private address. You cannot configure a data interface for management if the management interface is set to DHCP, because the default route, which must be **data-interfaces** (see the next bullet), might be overwritten with one received from the DHCP server.
- **Enter the IPv4 default gateway for the management interface**—If you want to use a data interface for Firewall Threat Defense access instead of the management interface, set the gateway to be **data-interfaces**. This setting forwards management traffic over the backplane so it can be routed through the FMC access data interface.
- **If your networking information has changed, you will need to reconnect**—If you are connected with SSH but you change the IP address at initial setup, you will be disconnected. Reconnect with the new IP address and password. Console connections are not affected.
- **Manage the device locally?**—Enter **YES** to configure the device for the device to be managed by either the Cloud-Delivered Firewall Management Center or Secure Firewall Device Manager.

  Manage the device locally?—Enter **NO** to configure the device for remote management with the on-premises Firewall Management Center.
- **Configure firewall mode?**—We recommend that you set the firewall mode at initial configuration. Changing the firewall mode after initial setup erases your running configuration. Note that data interface Firewall Threat Defense access is only supported in routed firewall mode.

Step 5 (Optional) Configure a data interface for Cloud-Delivered Firewall Management Center access.

configure network management-data-interface

You are then prompted to configure basic network settings for the data interface.

Note
You should use the console port when using this command. If you use SSH to the Management interface, you might get disconnected and have to reconnect to the console port. See below for more information about SSH usage.

See the following details for using this command. See About Data Interfaces, on page 3105 for more informatio.

- The original management interface cannot use DHCP if you want to use a data interface for management. If you did not set the IP address manually during initial setup, you can set it now using the **configure network** {**ipv4** | **ipv6**} **manual** command. If you did not already set the management interface gateway to **data-interfaces**, this command will set it now.
- When you onboard the device for Firewall Threat Defense management through Security Cloud Control, Security Cloud Control discovers and maintains the interface configuration, including the following settings: interface name and IP address, static route to the gateway, DNS servers, and DDNS server. For more information about the DNS server configuration, see below. You can later make changes to the access interface configuration, but make sure you don't make changes that can prevent the device or Security Cloud Control from re-establishing the management connection. If the management connection is disrupted, the device includes the **configure policy rollback** command to restore the previous deployment.
- This command sets the *data* interface DNS server. The Management DNS server that you set with the setup script (or using the **configure network dns servers** command) is used for management traffic. The data DNS server is used for DDNS (if configured) or for security policies applied to this interface.

  Also, local DNS servers are only retained if the DNS servers were discovered at initial registration. For example, if you registered the device using the Management interface, but then later configure a data interface using the **configure network management-data-interface** command, then you must manually configure all of these settings in Security Cloud Control, including the DNS servers, to match the device configuration.
- You can change the management interface after you onboard the Firewall Threat Defense for Firewall Threat Defense management through threat defense, to either the management interface or another data interface.
- The FQDN that you set in the setup wizard will be used for this interface.
- You can clear the entire device configuration as part of the command; you might use this option in a recovery scenario, but we do not suggest you use it for initial setup or normal operation.
- To disable data managemement, enter the **configure network management-data-interface disable** command.

Example:

```
> configure network management-data-interface
Data interface to use for management: ethernet1/1
Specify a name for the interface [outside]:
IP address (manual / dhcp) [dhcp]:
DDNS server update URL [none]:
https://jcrichton:pa$$w0rd17@domains.example.com/nic/update?hostname=<h>&myip=<a>
```

```
Do you wish to clear all the device configuration before applying ? (y/n) [n]:

Configuration done with option to allow FMC access from any network, if you wish to change
 the FMC access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.

>
```

Example:

```
> configure network management-data-interface
Data interface to use for management: ethernet1/1
Specify a name for the interface [outside]: internet
IP address (manual / dhcp) [dhcp]: manual
IPv4/IPv6 address: 10.10.6.7
Netmask/IPv6 Prefix: 255.255.255.0
Default Gateway: 10.10.6.1
Comma-separated list of DNS servers [none]: 208.67.222.222,208.67.220.220
DDNS server update URL [none]:
Do you wish to clear all the device configuration before applying ? (y/n) [n]:

Configuration done with option to allow FMC access from any network, if you wish to change
 the FMC access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.

>
```

Step 6 (Optional) Limit data interface access to Security Cloud Control on a specific network.

configure network management-data-interface client *ip_address netmask*

By default, all networks are allowed.

Delete Devices from Cloud-Delivered Firewall Management Center

Though devices may be registered to Cloud-Delivered Firewall Management Center, Security Cloud Control still manages device enrollment. You must delete the device from the Security Cloud Control dashboard to remove a device from Cloud-Delivered Firewall Management Center.

Note Security Cloud Control does not synchronize the deletion of devices that are associated with an AWS VPC environment. You must delete a device direclty from the AWS VPC UI. See AWS documentation for more information.

Procedure

Step 1 Log into Security Cloud Control and click **Security Devices**.

Step 2 Locate the device you want to delete by using the filters or search bar. Select it so the device row is highlighted. If your device is part of a high availability pair, locate and select the active device.

Step 3 In the Device Actions pane located to the right, click **Remove**.

Step 4 When prompted, select **OK** to confirm the removal of the selected device. Click **Cancel** to keep the device onboarded.

Troubleshooting

Use the following scenarios to troubleshoot any onboarding issues.

Troubleshoot Cloud-Delivered Firewall Management Center Connectivity with TCP

Use the following procedure to troubleshoot connectivity between the Cloud-Delivered Firewall Management Center and a Firewall Threat Defense device with TCP port 8305.

Procedure

Step 1 Log into Security Cloud Control.

Step 2 Navigate to **Tools & Services** in the left panel and select **Firewall Management Center** to open the **Services** page. Choose **Cloud-Delivered FMC** and locate the Cloud-Delivered Firewall Management Center's FQDN in the top righthand corner.

Step 3 Make sure the Firewall Threat Defense device's state in Security Cloud Control is currently **Onboarding**. Cloud-Delivered Firewall Management Center will not respond if the device is not in an onboarding state. If onboarding has failed, click **Retry Onboarding**.

Step 4 Log into the Firewall Threat Defense device with SSH.

Step 5 Enter into Expert mode with the following command:

```
> expert
admin@devicename:~$
```

Step 6 Execute a TCP handshake:

```
admin@devicename:~$ nc -v xxxxxx.cdo.cisco.com 8305
Connection to xxxxxx.cdo.cisco.com 8305 port [tcp/*] succeeded!
^C (CTRL-C to exit netcat)
admin@devicename:~$.
```

If there is still no response from the Cloud-Delivered Firewall Management Center, then there is a chance that outbound port TCP 8305 may be blocked upstream from your Firewall Threat Defense device and that

network path will need to be assured before your Firewall Threat Defense will be able to connect to Cloud-Delivered Firewall Management Center.

Troubleshoot Firewall Threat Defense Device Connectivity

Test connectivity to the Internet from the management-plane of the Firewall Threat Defense device:

Procedure

	Command or Action	Purpose
Step 1	Log into the Firewall Threat Defense device with SSH.	

Step 2	Ping one or both of the following:	`ping system 208.67.222.222`
		`ping system cisco.com`
		If either of these tests fail, there is likely an L1-L3 issue and you will need to check your management networking configuration (`show network`) and/or a DNS issue.

Troubleshoot Device Connectivity Loss After Cloud-Delivered Firewall Management Center Update

A Cloud-Delivered Firewall Management Center is assigned a dynamic IP address when it is added to a Security Cloud Control tenant. When the management center is updated, the management center receives a new dynamic IP address.

If you have a firewall inspecting the outbound traffic from your threat defense device to the Cloud-Delivered Firewall Management Center, your firewall rules must allow the threat defense traffic to flow to the FQDN and port of the management center rather than its IP address, or the management center will not be able to manage your threat defense device.

For example, if your network traffic rule allowing management traffic from your threat defense device to the Cloud-Delivered Firewall Management Center looks like this:

```
allow all traffic <my-threat-defense-ip-src> to 200.165.200.225
```

where 200.165.200.225 is the management address of the Cloud-Delivered Firewall Management Center, change the one allow rule to these two allow rules as both ports 443 and 8305 need to be open:

```
allow all traffic <my-threat-defense-ip-src > to <my-cdfFMC-FQDN>:443
```

```
allow all traffic <my-threat-defense-ip-src > to <my-cdfFMC-FQDN>:8305
```

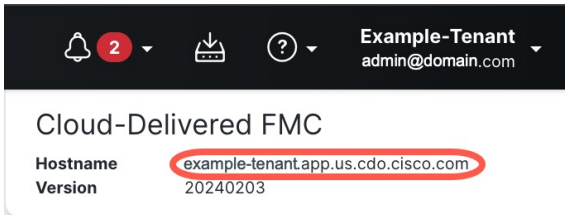
See “Network Requirements” in Prerequisites to Onboard a Device to Cloud-delivered Firewall Management Center for more port information.

Where do I find the domain name of my Cloud-Delivered Firewall Management Center?

Where do I find the domain name of my Cloud-Delivered Firewall Management Center?

1. Log in to Security Cloud Control.
2. In the left pane, click **Administration** > **Integrations** > **Firewall Management Center**
3. Select **Cloud-Delivered FMC** in the FMC table.
4. In the top-right corner of the screen, you will see the Hostname of the management center. This is the FQDN.

Figure 7: Cloud-delivered FMC FQDN



Troubleshoot Onboarding a Device to the Cloud-Delivered Firewall Management Center Using the CLI Registration Key

Error: Device Remains in Pending Setup State After Onboarding

When a device fails to register, the device's connectivity status is displayed as **Pending Setup**. In the panel located to the right, Security Cloud Control displays a **Registration Failed** message as well as a **Retry Onboarding** button to immediately allow you to reattempt onboarding the device.

If you fail to execute the `configuration manager` command in the device CLI within 3 mins after onboarding it to Security Cloud Control, the device's registration attempt expires and results in a registrationg failure. Use the following procedure to resolve the issue:

Procedure

Step 1 Log into Security Cloud Control and navigate to the **Security Devices** page. Locate the device that failed to register.

Step 2 In the panel located to the right, locate the **Registration Failed** window. Beside the device's CLI registration key, click **Copy**. This actions copies the CLI key to a local clipboard.

Step 3 Open an SSH connection to the device and log in as `admin`.

Step 4 Paste the CLI registration key into the device's CLI interface. In the CLI, enter **Y** to complete the registration. If your device was previously managed by Firewall Device Manager, enter **Yes** to confirm the submission.

Troubleshoot Onboarding a Device to Cloud-Delivered Firewall Management Center Using the Serial Number

Device is Offline or Unreachable

If the device is unreachable during the onboarding proess, or at any point post-onboarding, Security Cloud Control displays an **Unreachable** connectivity status. The device will not be able to fully onboard to Security Cloud Control until the device is able to connect. The following scenarios might be the cause:

- The device is cabled incorrectly.
- Your network may require a static IP address for the device.
- Your network uses custom DNS, or there is external DNS blocking the network.
- If your device is associated with the European region (https://defenseorchestrator.eu/), you may need to enable PPPoE authentication.For other domains, review the domain requirements.
- The device may be blocked by a firewall, or is incorrectly blocking a port for connectivity. Review the device Network Requirements, on page 22 and confirm the correct outgoing ports are enabled.

Error: Serial Number Already Claimed

The Device was Purchased From an External Vendor

If the device was purchased from an external vendor and fails to onboard with a **Serial Number Already Claimed** error, it's possible the device is still associated to the vendor's tenant. Use the following steps to claim the device and its serial number:

1. Delete the device from your Security Cloud Control tenant.
2. Install the FXOS image on the device. For more information, see the "Reimage Procedures" chapter of the Cisco FXOS Troubleshooting Guide for the Firepower 1000/21000 and Secure Firewall 3100 Firepower Threat Defense guide.
3. Connect a laptop to the device's console port.
4. Connect to the FXOS CLI and log in as **admin**.
5. In the FXOS CLI, connect to **local-mgmt** with the `firepower # connect local-mgmt` command.
6. Execute the `firepower(local-mgmt) # cloud deregister` command to deregister the device from the cloud tenancy.
7. Once the device is successfully unregistered, the CLI interface returns a success message. An example of the message:

```
Example: firepower(local-mgmt) # cloud deregister Release Image Detected RESULT=success
 MESSAGE=SUCCESS 10, X-Flow-Id: 2b3c9e8b-76c3-4764-91e4-cfd9828e73f9
```

Note If the device was never registered to another Security Cloud Control tenant, the message above states `RESULT=success MESSAGE=DEVICE_NOT_FOUND`.

8. Onboard the device to your Security Cloud Control tenant with its serial number. See Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning, on page 24 for more information.

The Device is Claimed By a Security Cloud Control Tenant in Another Region

The device may have been previously managed by another Security Cloud Control instance in a different region and is still registered to that tenant.

If you **do** have access to the tenant the device is currently registered to, use the following procedure:

1. Delete the device from the incorrect Security Cloud Control tenant.
2. Log into the device's Firewall Device Manager UI.
3. Navigate to **System Settings** > **Cloud Services**.
4. Click **Cloud Services** and select **Unregister Cloud Services** from the drop-down list.
5. Confirm the action and click **Unregister**. This action generates a warning to indicate that the device has been removed from Security Cloud Control. This is expected behavior.
6. Log into Security Cloud Control tenant in the correct region and onboard the device. See Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning, on page 24 for more information.
7. Navigate to **System Settings** > **Cloud Services**.
8. Click **Cloud Services** and select **Unregister Cloud Services** from the drop-down list.
9. Select the **Auto-enroll with Tenancy from Security Cloud Control** and click **Register**. The device maps to the new tenant that belongs to the new region and Security Cloud Control onboards the device.

If you **do not** have access to the tenant, use the procedure below:

1. Connect to the FXOS CLI from the console port and log in as **admin**. For information on how to log into the FXOS CLI, see Accessing the FXOS CLI.
2. In the FXOS CLI, connect to **local-mgmt** with the `firepower # connect local-mgmt` command.
3. Execute the `firepower(local-mgmt) # cloud deregister` command to deregister the device from the cloud tenancy.
4. Once the device is successfully unregistered, the CLI interface returns a success message. An example of the message:

```
Example: firepower(local-mgmt) # cloud deregister Release Image Detected RESULT=success
 MESSAGE=SUCCESS 10, X-Flow-Id: 2b3c9e8b-76c3-4764-91e4-cfd9828e73f9
```



Note If the device was never registered to another Security Cloud Control tenant, the message above states `RESULT=success MESSAGE=DEVICE_NOT_FOUND`.

5. In your Security Cloud Control tenant in the correct domain, onboard the device. See Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning, on page 24 for more information.

6. In the device's Firewall Device Manager UI, navigate to **System Settings** > **Cloud Services**.
7. Select the **Auto-enroll with Tenancy from Security Cloud Control** and click **Register**. The device maps to the new tenant that belongs to the new region and Security Cloud Control onboards the device.

Error: Claim Error

If you enter the wrong serial number when onboarding a device, Security Cloud Control generates a **Claim Error** status.



Note Confirm that the device is claimed in the correct region within Security Cloud Control.

Resolve this issue with the provedure below:

Procedure

Step 1 Log into Security Cloud Control and navigate to the **Security Devices** page. Locate the device with the error.

Step 2 Select the device so it is highlighted and **Remove** the device from Security Cloud Control.

Step 3 Confirm the following:

- The device is online and can reach the internet.
- The device has not already been onboarded to your Security Cloud Control instance or claimed by a Security Cloud Control tenant in another region.

Step 4 Locate the device's serial number. You can use one of the following methods:

- For the 1000, 2100 and 3100 series model, locate the serial number on the physical device.
- Open an SSH connection to the device and issue the `show serial-number` command.

Step 5 In Security Cloud Control, onboard the device with the correct serial number. See Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning, on page 24 for more information.

Error: Failed to Claim

If you see an **Error: Failed to Claim** connectivity status or error message after attempting to onboard a device, the following might be the cause:

- The Security Services Exchange platform may have temporary issues that result in no connectivity.
- The Security Cloud Control server may be down.

Follow the procedure below to resolve this issue:

Procedure

Step 1 Log into Security Cloud Control and navigate to the **Security Devices** page. Locate the device that failed to register.

Step 2 Select the device so it is highlighted and **Remove** the device from your Security Cloud Control tenant.

Step 3 Wait at least 10 minutes before attempting to onboard the device back to your Security Cloud Control tenant. See Onboard a Firewall Threat Defense Device to Cloud-Delivered Firewall Management Center using Zero-Touch Provisioning, on page 24 for more information.

If you are still unable to claim the device, review the device's workflow to see is there is an error message. If there is,Export the Workflow and open a support case to further troubleshoot the issue.

Error: Provisional Error

Device Password Has Not Been Changed

If you did not change the default password of the device when configuring the device for remote management and selected the **No, this device has been logged into and configured for a manager** option when onboarding the device to Security Cloud Control, the device will generate an **UnProvisioned** connectivity status in the **Security Devices** page.

Use the following procedure to resolve this issue:

1. Log into Security Cloud Control and navigate to the **Security Devices** page.
2. Locate and select the device with the **UnProvisioned** connectivity status so it is highlighted.
3. In the pane located to the right, locate the **Change Password** window.
4. Click **Change Password** and enter a new password for your device. This overwrites the default password.

It may take a few minutes for the device to onboard and fully synchronize to Security Cloud Control.

Device Password Has Already Been Changed

If you **did** change the default password of the device when configuring the device for remote management and selected the **Is this a new device that has never been logged into or configured before?** option when onboarding the device to Security Cloud Control, Security Cloud Control generates an **UnProvisioned** connectivity status in the **Security Devices** page.

Use the following procedure to resolve this issue:

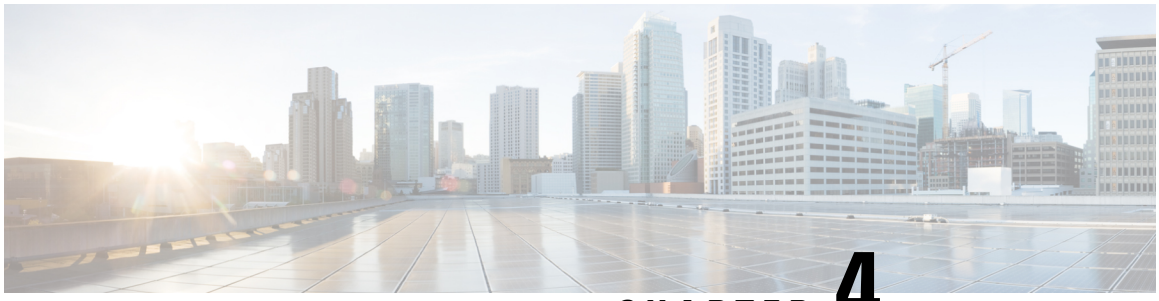
1. Log into Security Cloud Control and navigate to the **Security Devices** page.
2. Locate and select the device with the **UnProvisioned** connectivity status so it is highlighted.
3. In the pane located to the right, locate the **Confirm and Proceed** window.
4. Click **Confirm and Proceed**. This action ignores the password that was provided in the onboarding wizard and reinstates the default password for the device. Security Cloud Control then continues to onboard the device.

Other Provisional Error Scenarios

Regardless of the default password configuration of the device, it is still possible for a device to result in an **UnProvisioned** conectivity status during the onboarding process. If you confirm the password selection in the onboarding wizard is accurate for the state of the device, consider the following options to resolve the issue:

- Select the device so it is highlighted. In the window located on the right pane of the screen, click **Retry** to force Security Cloud Control to re-onboard the device with existing provisional parameters.
- Delete the device from the **Security Devices** page and attempt to re-onboard the device.
- In the device's Firewall Device Manager UI, navigate to **System Settings** > **Cloud Services**. Select the **Auto-enroll with Tenancy from Security Cloud Control** and click **Register**.

If you are still unable to claim the device, review the device's workflow to see is there is an error message. If there is,Export the Workflow and open a support case to further troubleshoot the issue.



CHAPTER 4

Migrate On-Premises Firewall Management Center Managed Secure Firewall Threat Defense to Cloud-delivered Firewall Management Center

- Overview of Firewall Threat Defense to Cloud-Delivered Firewall Management Center migration, on page 63
- End-to-End Workflow for Migrating Firewall Threat Defense to Cloud-Delivered Firewall Management Center, on page 66
- Event and Analytics Management options during Firewall Threat Defense migration, on page 67
- Supported features, on page 68
- Unsupported features, on page 70
- User Identity Migration Guidelines and Limitations for Firewall Threat Defense Devices, on page 71
- Migration guidelines and limitations for VPN configuration, on page 71
- Supported On-Premises Firewall Management Center and Firewall Threat Defense software versions for migration, on page 73
- Prerequisites for Migrating Firewall Threat Defense to Cloud-Delivered Firewall Management Center, on page 75
- Migrate Firewall Threat Defense to Cloud-Delivered Firewall Management Center, on page 78
- View a Firewall Threat Defense Migration Job, on page 81
- Enable Notification Settings, on page 86
- Verify Firewall Threat Defense Connectivity with Cloud-Delivered Firewall Management Center, on page 87
- Use the Troubleshoot Feature for Failed Migrations, on page 88

Overview of Firewall Threat Defense to Cloud-Delivered Firewall Management Center migration

Introduction

This feature enables you to migrate Firewall Threat Defense devices from On-Premises Firewall Management Center to Cloud-Delivered Firewall Management Center through Security Cloud Control Firewall Management. Migration enables organizations to move device management from on-premises infrastructure to Cisco's cloud-based management. This migration preserves existing policies, objects, and configurations.

Customer benefits

- Preserves existing policies and configuration
- Imports device-specific configuration, including interfaces, routing, and related settings
- Handles duplicate policy and object names according to defined conflict-resolution rules
- Transfers feature licenses automatically
- Supports high availability pairs and clusters
- Provides 14-day evaluation window before permanent commit
- Allows analytics to remain on-prem or move to cloud
- Provides an intuitive troubleshooting page with status summaries and step-by-step progress to pinpoint issues.
- Provides migration reports for validation and audit

User role changes

The user roles of the On-Premises Firewall Management Center are no longer applicable in Security Cloud Control Firewall Management after migration. Your authorization to perform tasks on the migrated device is based on your user role in Security Cloud Control Firewall Management. See the Users topic to understand the On-Premises Firewall Management Center Center and Cloud-Delivered Firewall Management Center user role mapping.

How Firewall Threat Defense Migration to Cloud-Delivered Firewall Management Center Works

Security Cloud Control Firewall Management performs migration in the following sequence:

1. Onboards on-premises Firewall Management Center to Security Cloud Control Firewall Management.
2. Exports configurations from on-premises Firewall Management Center.
3. Imports shared policies, objects, and device-specific configurations to Cloud-Delivered Firewall Management Center.
4. Registers Firewall Threat Defense devices with Cloud-Delivered Firewall Management Center.
5. Starts a 14-day evaluation window to commit or revert changes.
6. Allows a commit or revert changes.

After migration:

- Cloud-Delivered Firewall Management Center manages devices and their configuration.
- Events and analytics can be handled by either the on-premises Firewall Management Center (analytics-only mode) or Security Cloud Control Firewall Management.
- Devices appear on the **Security Devices** page.

About 14-day evaluation period

When a migration job is successful, you have 14 days to test and assess migration changes using Cloud-Delivered Firewall Management Center. If you are convinced about the migration changes, we recommend that you commit the devices manually, and not wait for Secure Firewall Management Center to automatically commit the migration changes. You have a 14-day evaluation period to review and assess the migration changes that are made to the devices before Secure Firewall Management Center automatically commits them.

After successful migration:

- You have 14 days to evaluate configuration.
- You can undo the changes and continue managing the device with On-Premises Firewall Management Center.
- You can commit migration changes to Cloud-Delivered Firewall Management Center.

After 14 days evaluation period:

- Changes are automatically committed.
- Revert is no longer possible after auto-commit.

During 14 day evaluation window:

- High availability break may not be supported (depending on onboarding method).
- Performing advanced high availability or cluster operations may result in migration commit failure.
- Deleting devices is not allowed.
- Changes made in Cloud-Delivered Firewall Management Center are lost, if reverted.
- Do not create or modify domains or subdomains on the On-Premises Firewall Management Center if the domain change includes a Firewall Threat Defense device that was migrated to Cloud-Delivered Firewall Management Center. Commit the migration changes in Cloud-Delivered Firewall Management Center, or revert device management to the on-premises On-Premises Firewall Management Center, before you make domain changes that include the migrated device.

  If you create or modify a domain before the migration job is committed, deployment of the device and domain configuration can fail. You can continue to create or modify domains on the On-Premises Firewall Management Center for devices that remain managed by that On-Premises Firewall Management Center and are not part of an uncommitted migration job.

How Firewall Threat Defense licenses are handled during migration to Cloud-Delivered Firewall Management Center

- When the Firewall Threat Defense is migrated to the cloud, all feature licenses associated with the device are transferred to Security Cloud Control and released from the Firewall Management Center to the Smart License pool. The device reclaims the device-specific licenses during its registration with Security Cloud Control. You need not apply license on the device again.
- The device-specific licenses are not required if you want to keep devices in the Firewall Management Center for analytics.

End-to-End Workflow for Migrating Firewall Threat Defense to Cloud-Delivered Firewall Management Center

The following flowchart illustrates the high-level workflow for migrating Firewall Threat Defense to Cloud-Delivered Firewall Management Center.

Event and Analytics Management options during Firewall Threat Defense migration

The events and analytics management can be retained in the on-premises Firewall Management Center or transferred to Security Cloud Control Firewall Management, where the devices must be configured to send events to Security Cloud Control Firewall Management.

While initiating the migration process, you are allowed to choose the manager where the device events must be sent for analytics:

- Retain analytics on the on-premises Firewall Management Center, or
- Transfer analytics to Security Cloud Control Firewall Management.

Special requirement for FMC 1000/2500/4500

If you are migrating from an on-premises Firewall Management Center 1000/2500/4500, you **cannot** retain analytics on the on-premises Firewall Management Center due to limited availability.

In this case, devices must send events to:

- Security Analytics and Logging (On-Prem), or
- Security Analytics and Logging (SaaS).

See Cisco Security Analytics and Logging.

If you retain on-premises Firewall Management Center for analytics

- Security Cloud Control Firewall Management becomes the configuration manager.
- The devices remain on the on-premises Firewall Management Center in **analytics-only mode**.
- Devices continue sending events to the on-premises Firewall Management Center.
- Security Cloud Control Firewall Management manages configuration changes only.

Note The maximum number of supported Secure Firewall Threat Defense devices in Cloud-Delivered Firewall Management Center includes devices that are deployed for analytics on Firewall Management Center and onboarded to Cloud-Delivered Firewall Management Center for management. These devices count toward the platform device capacity limit and must be included in your capacity planning to ensure proper sizing and performance.

If you select Security Cloud Control Firewall Management for analytics

- Secure Firewall Management Center becomes both:
  - Configuration manager, and
  - Analytics manager.

- The devices are removed from the on-premises Firewall Management Center.
- You must configure the devices to send events to the Cisco cloud.
- Events can be sent using:
  - Security Services Exchange (SSE), or
  - Secure Event Connector (SEC)
- Events are processed by Cisco Secure Analytics and Logging (SAL) in the cloud.



Note If you initially select on-premises Firewall Management Center **for analytics**, you have a 14-day evaluation period during which you can change the analytics destination to Secure Firewall Management Center.

After either:

- Manually committing the migration, or
- The 14-day evaluation period expires, the analytics selection becomes permanent. Events will continue to be sent to the selected destination, and the setting cannot be changed.

To modify analytics settings after the evaluation period or after commit, refer to the migration troubleshooting procedures. see Resolving Firewall Threat Defense Migration to Cloud-Delivered Firewall Management Center Issues, on page 90

eStreamer Server Streaming

When you manage a Firewall Threat Defense device with Cloud-Delivered Firewall Management Center, the device supports sending only fully-qualified events (FQE) to eStreamer clients. If you have configured eStreamer clients in the on-premises Firewall Management Center, ensure that the clients support the detailed data formats used by FQE when you migrate the device management to Cloud-Delivered Firewall Management Center. Any legacy clients, security information and event management (SIEM) systems, or log management solutions that do not support the data format of FQE or lack the necessary storage to handle the larger volume of FQE data will not work when you migrate.

Supported features

Shared policies imported

The following policies are imported first and then followed by the device configuration:

- Access control
- IPS
- SSL
- Prefilter
- NAT

- QoS
- Identity
- Platform settings
- Flex config
- Network analysis
- DNS
- Malware & file
- Health
- Remote Access VPN
- Site-to-Site VPN

Handling duplicate names

If a policy or object name in Security Cloud Control Firewall Management matches one imported from the On-Premises Firewall Management Center, Security Cloud Control Firewall Management will take the following actions after a successful management change.

Policies, Objects	Condition	Action
Access control, SSL, IPS, Prefilter, NAT, QoS, Identity, Platform settings, Network analysis, DNS, Malware & File policies.	Name of the Cloud-Delivered Firewall Management Center policy matches the on-premises Firewall Management Center policy.	Cloud-Delivered Firewall Management Center policy is used.
RA VPN Default group policy **DfltGrpPolicy**	The default group policy**DfltGrpPolicy** from the on-premises Firewall Management Center is ignored.	Existing Cloud-Delivered Firewall Management Center default group policy **DfltGrpPolicy** is used.
Network, Port objects	Name and content of network and port objects in the Cloud-Delivered Firewall Management Center match the ones in the on-premises Firewall Management Center.	Existing Cloud-Delivered Firewall Management Center network and port objects with the same name and content are used. If the object has the same name but different content, an object override is created.
All other objects		Existing Cloud-Delivered Firewall Management Center object is used.

Note Syslog alert objects used in the access control policy are imported.

Unsupported features

Migration of a Firewall Threat Defense device registered only for analytics-only with the Firewall Management Center feature is not currently supported.

The following configuration are not imported from the Firewall Management Center to Security Cloud Control as part of migration:

- Custom Widgets, Application Detectors, Correlation, SNMP and Email Alerts, Scanners, Groups, Dynamic Access Policy, Custom AMP Configuration, Users, Domains, Scheduled Deployment Tasks, ISE configuration, Scheduled GeoDB Updates, Threat Intelligence Director configuration, Dynamic Analysis Connections.
- ISE internal certificate object is not imported as part of the migration. You must export a new system certificate or a certificate and its associated private key from ISE and import it into Security Cloud Control Firewall Management.

Secure firewall recommended rules

Migrating Firewall Threat Defense to the cloud migrates existing rule recommendations linked to intrusion policies. However, the Cloud-Delivered Firewall Management Center does not generate new rule recommendations or auto-update migrated ones post-migration, as it does not support rule recommendations. See Auto Cisco Recommended Rules.

Remove custom network analysis policies

Before migration, Custom Network Analysis policies must be removed from the On-Premises Firewall Management Center.

1. Log on to the on-premises Firewall Management Center.
2. Choose **Policies** > **Access Control**.
3. Click the edit icon on the access control policy you want to disassociate the custom NAP and then click the **Advanced** tab.
4. In the **Network Analysis and Intrusion Policies** area, click the edit icon.
5. In the **Default Network Analysis Policy** list, select a system-provided policy.
6. Click **OK**.
7. Click **Save** to save the changes and then click **Deploy** to download the changes to the device.

After migration, you can manually create the Network Analysis Policy in Security Cloud Control Firewall Management.

User Identity Migration Guidelines and Limitations for Firewall Threat Defense Devices

Before you migrate an on-premises Firewall Threat Defense to Cloud-Delivered Firewall Management Center, you must prepare and also deploy as soon as possible after the migration if any access control policies reference identity objects discussed in this topic. To confirm, click **Policies** > **Access Control heading** > **Access Control** and examine your access control policies and rules.

If none of your access control policies reference identity objects (in particular, users and groups), you can ignore these guidelines.

Before migrating

Before migrating, on the On-Prem Firewall Management Center, click **Integration** > **Other Integrations** > **Identity Sources** and see if you have any Cisco ISE/ISE-PIC or Passive Identity Agent identity sources defined.

- If you have Cisco ISE/ISE-PIC or Passive Identity Agent identity sources defined, create the Cisco ISE/ISE-PIC identity sources on Cloud-Delivered Firewall Management Center then migrate the device as discussed in the following paragraphs.
- If no Cisco ISE/ISE-PIC or Passive Identity Agent identity sources are defined, migrate the device as discussed in the following paragraphs.

Migrate the device

Migration is discussed in Overview of Firewall Threat Defense to Cloud-Delivered Firewall Management Center migration, on page 63. To avoid traffic disruption, when you migrate the device, we strongly recommend you either:

- Check the **Auto deploy to FTDs after successful migration** check box.
- Deploy policies immediately after migration is complete.

Migration guidelines and limitations for VPN configuration

Keep the following in mind when you migrate a device with VPN configuration.

Migration guidelines for Remote Access VPN

Security Cloud Control Firewall Management imports all remote access VPN policy settings, with the following exceptions:

- Object overrides are not migrated.

  If overrides are used in the address pool object, manually add them to the imported object in Security Cloud Control Firewall Management after migration.

- Local users.

If the authentication server uses a local database, the associated local realm object is imported into Security Cloud Control Firewall Management. However, you must manually add the local users to the imported local realm object after migration. See Create a Realm and Realm Directory.

- Load-balancing configuration is not migrated.
- Certificate enrollment with domain configuration.

  Perform the following after migration to enroll the certificate with domain configuration:

  1. Choose **Security Devices**.
  2. Select the migrated FTD and in the **Device Management** on the right, click **Device Overview**.
  3. Choose **Devices** > **Certificates**.

     Perform one of the following tasks:

     - If the certificates are imported in an **Error** state, click the **Refresh certificate status** icon to synchronize the certificate status with the device. The certificate status turns green.
     - If the certificates are not imported, you must manually add the certificates defined in the Remote Access VPN policy that is configured in the Firewall Management Center.

Migration guidelines for Site-to-Site VPN

- After selecting a Firewall Threat Defense device, all peers are automatically selected from various topologies, as all devices within a topology must migrate together.
- Migration wizard doesn't list the extranet devices that are associated with them, but they will still be automatically included during the migration process.

All settings will be migrated, with the following exceptions:

- Network object overrides are not migrated. They must be manually added to the imported object in Security Cloud Control Firewall Management after migration.
- If the authentication type configured as **Preshared Automatic Key** in the On-Premises Firewall Management Center:
  - Cloud-Delivered Firewall Management Center generates new pre-shared key post-deployment.
  - Existing tunnels remain active.
  - New tunnels use newly generated key.
- If devices are moved to Security Cloud Control Firewall Management but the changes are not yet committed, the associated site-to-site VPN policy can still be modified in the On-Premises Firewall Management Center, however, it those changes do not update the device configuration in Security Cloud Control Firewall Management.
- Devices configured for SASE tunnels on Cisco Umbrella are not migrated.

Supported On-Premises Firewall Management Center and Firewall Threat Defense software versions for migration

Minimum supported versions

- Minimum On-Premises Firewall Management Center: 7.2
- Minimum Firewall Threat Defense: 7.2.x (not supported for Version 7.1)

Physical On-Premises Firewall Management Center 1000/2500/4500 Model-Managed Firewall Threat Defense

- Migration is supported from **physical FMC 1000/2500/4500** models.
- *Temporary upgrade from 7.0 to 7.4* is supported for migration purposes only. The On-Premises Firewall Management Center Version 7.4 is "unsupported for general operations," but allowed as an interim step until migration completes. You can download the upgrade package here.
- After migration, do the following:
  1. Return to On-Premises Firewall Management Center to a supported version after migration.
  2. Remove re-migrated devices.
  3. Reimage back to 7.2.x.
  4. Re-register the devices.

Note Unzip (but do not untar) the upgrade package before uploading it to the On-Premises Firewall Management Center. To upgrade to Version 7.4, see Cisco Secure Firewall Management Center Upgrade Guide, Version 6.0-7.0.

We recommend upgrading the devices to Version 7.2.x before upgrading the On-Premises Firewall Management Center to Version 7.4.

Important An upgrade is necessary because Version 7.0 On-Premises Firewall Management Centers do not support cloud device migration. Only Version 7.4 is supported during the migration and evaluation process. These On-Premises Firewall Management Centers cannot run intermediate versions. Only standalone and high availability Firewall Threat Defense devices running Version 7.0.3+ (7.0.5 recommended) are eligible for migration.

Firewall Threat Defense in a High Availability Pair

- High availability pairs migrate as a single node.
- Both active and standby Firewall Threat Defense devices migrate together.

Important We highly recommend committing manager changes before performing advanced operations, like creating or breaking high availability configurations from the On-Premises Firewall Management Center on the devices being migrated. Carrying out these tasks during the evaluation period is unsupported and may lead to migration commit failure.

Firewall Management Center in a High Availability Pair

- Onboard the active On-Premises Firewall Management Center and not the standby.
- Use either the auto-onboarding method or the credentials method for onboarding the On-Premises Firewall Management Center.

Note If you have already onboarded a standalone On-Premises Firewall Management Center and later configured it as a standby, delete the standby On-Premises Firewall Management Center and onboard the active one.

Points to Remember:

- **Auto-onboarding On-Premises Firewall Management Center Method**
  - High availability break is not supported during the 14 days evaluation period.

    You can break high availability after changes are committed, manually or automatically, after the evaluation period.
  - High availability switchover is supported during the 14 days evaluation period.
- **Onboarding On-Premises Firewall Management Center Method Using Credentials**
  - High availability break and high availability switchover are not supported during the 14 days evaluation period. You can perform these operations after committing the changes manually or automatically after the evaluation period.
  - After a switchover, onboard the new active unit, which was previously in standby mode, and then start a migration job on the devices.

Note During the 14-day evaluation period after migrating a Firewall Threat Defense device, if you have performed advanced operations on the device, such as converting it to a multi-instance chassis, note that the evaluation period becomes nullified, and you will not be able to revert the device management back to the On-Premises Firewall Management Center.

If you want to manage a multi-instance chassis using Cloud-Delivered Firewall Management Center, onboard the chassis manually to Security Cloud Control.

Firewall Threat Defense cluster

Firewall Threat Defense cluster migration is supported on these platforms with minimum versions:

Platforms	Minimum Firewall Threat Defense	Minimum On-Premises Firewall Management Center Version
VMware, KVM	7.2.1	7.4.1
AWS, GCP	7.2.1	7.4.1
Azure	7.3	7.4.1
Secure Firewall 3100	7.2.1	7.4.1
Firepower 4100	7.2.x	7.4.1
Secure Firewall 4200	7.4	7.4.1
Firepower 9300	7.2.x	7.4.1



Important Before migrating the Firewall Threat Defense cluster, remember the following points:

- Do not migrate the Firewall Threat Defense cluster during clustering-related operations from the On-Premises Firewall Management Center.
- After migration, manually commit manager changes before performing advanced operations, such as adding a node, breaking a node, breaking a cluster, or converting a device from native mode to multi-instance (container) mode from the on-premises Firewall Management Center. Performing such tasks during the evaluation period may result in migration commit failure.

Migration Support for Multi-instance Firewall Threat Defense Devices

You can migrate a multi-instance Firewall Threat Defense device that is part of a chassis (Secure Firewall 3100/4200). The multi-instance device gets listed as one of the devices in the **Select Devices** page, and you can choose it to proceed with the migration. This migration is supported in Firewall Threat Defense device versions 7.6 or later.

If you decide to do such a multi-instance Firewall Threat Defense device migration, it is strongly recommended that you unregister the corresponding chassis device manually from the on-premises Firewall Management Center and onboard it to Security Cloud Control, by navigating to the **Security Devices** page.

Prerequisites for Migrating Firewall Threat Defense to Cloud-Delivered Firewall Management Center

Before you begin the process, ensure that the following prerequisites are met:

- **DNS Server Configuration**:

  The Firewall Threat Defenses must have the correct DNS server configuration to resolve Cloud-Delivered Firewall Management Center hostnames. For details about checking connectivity with Cloud-Delivered

Firewall Management Center, see Check device connectivity with Cloud-Delivered Firewall Management Center.

- **Network Access**:

  The required network access must be enabled for Firewall Threat Defenses to reach Cloud-Delivered Firewall Management Center through the TCP port 8305. Note that outbound connectivity from the Firewall Threat Defenses to Cloud-Delivered Firewall Management Center is sufficient.

- **Firewall Threat Defense Outbound Port 443**:

  The Firewall Threat Defenses must be able to access port 443 and 8305 of Cloud-Delivered Firewall Management Center.

- **On-Premises Firewall Management Center Outbound Port 443**:

  The on-premises Firewall Management Center must have access to outbound port 443 open to reach "*.cdo.cisco.com" domain.

- The on-premises Firewall Management Center is onboarded to Security Cloud Control. Onboarding the on-premises Firewall Management Center also onboards all the Firewall Threat Defense devices registered to that on-premises Firewall Management Center. See Onboard an On-Prem FMC.

Note Create a new user in the on-premises Firewall Management Center with Administrator role or a custom user role with "Devices" and "System" permissions for onboarding purposes.

Caution If you onboard an on-premises Firewall Management Center to Security Cloud Control and simultaneously sign in to that on-premises Firewall Management Center with the same user name, the onboarding fails.

- For the on-premises Firewall Management Center 1000/2500/4500 migration:
  - Run Version 7.4 (available for these models on a temporary basis). We recommend devices be running Version 7.0.5.
  - We recommend that you create a backup of on-premises Firewall Management Center.

    For versions on-premises Firewall Management Center Version 6.5 to 7.1, see the *Back up the FMC* topic in the Firepower Management Center Configuration Guide.

    For on-premises Firewall Management Center Version 7.2 and later, see the *Back up the Management Center* topic in the Cisco Secure Firewall Management Center Administration Guide.

- The Firewall Threat Defense devices must be synchronized and not have pending changes on them. The migration fails on a device if Security Cloud Control identifies pending changes on that device.
- All peer devices in a site-to-site VPN topology must be online and have no pending deployment.
- On-Premises Firewall Management Center should allow outbound HTTP/HTTPS to upload configurations to Amazon S3.

- Security Cloud Control imports Syslog alert object used in the access control policy from the on-premises Firewall Management Center. If Security Cloud Control already contains an alert object with the same name but a different type (SNMP, Email), it is reused during configuration import.

  The user must check whether the Syslog object name matches the existing SNMP or Email alert object in Security Cloud Control. If the name matches, you must rename the Syslog object in the on-premises Firewall Management Center before starting the migration process.

- If you attempt to migrate firewalls with modified system defined FlexConfig text objects from an on-premises Firewall Management Center to the Cloud-Delivered Firewall Management Center, the values of the modified system defined FlexConfig text objects are not migrated to the Cloud-Delivered Firewall Management Center, and the deployment will fail.

  To avoid this, perform these tasks before you start the migration:

  - Copy the modified system defined FlexConfig text object values from the on-premises Firewall Management Center to Cloud-Delivered Firewall Management Center before migration.
  - Initiate migration from on-premises Firewall Management Center to Cloud-Delivered Firewall Management Center after verifying the predefined FlexConfig text objects.

Points to remember

Policy and Object Name Conflicts:

If a policy in Cloud-Delivered Firewall Management Center shares the same name as one in on-premises Firewall Management Center (excluding network, port, URL, and VLAN objects), the Cloud-Delivered Firewall Management Centerversion will be retained. See Handling Shared Policies and Objects.

Policy and Object Updates:

After migration, policies and objects aren't auto-synced between on-premises Firewall Management Center and Cloud-Delivered Firewall Management Center. For phased migrations, you must update them manually.

Evaluation Period:

You can revert to on-premises Firewall Management Center within 14 days of migration. After that, changes are finalized automatically. See Revert the Threat Defense Management to On-Prem Firewall Management Center.

High availability failover link must be up

The high availability failover link should be up for a successful migration. Before initiating the migration process on Security Cloud Control, determine the health status of the failover link on the on-premises Firewall Management Center.

1. Identify the failover interfaces of all HA pairs you want to migrate to Cloud-Delivered Firewall Management Center.
   a. Choose **Devices** > **Device Management**.
   b. Next to the device high-availability pair you want to edit, click **Edit** ( ).
   c. Click the **High Availability** tab.
   d. In the **High Availability Link** area, the **Interface** field shows the failover interface used in the pair.
   e. Identify the interfaces used for failover communication if there are multiple HA pairs for migration.

2. Check the health status of the failover interfaces.
   a. Choose **Devices** > **Device Management**.
   b. Next to the device high-availability pair you want, click **Health Monitor**.
   c. In the left pane, expand the high availability pair to see the Firewall Threat Defense devices.
   d. Click the device indicated in the exclamation mark ( ).
   e. Click the **Critical** button at the top.

      The **Interface Status** shows the errors associated with interfaces.
   f. If the failover interface is down, the **Interface 'failover_interfacename' has no link** message is displayed.

Note However, you can migrate the HA pair to Cloud-Delivered Firewall Management Center if you see any other data interface issues except for the failover interface.

   g. Rectify the issue and click **Sync from onprem fmc now** to obtain the latest changes on the device.

Migrate Firewall Threat Defense to Cloud-Delivered Firewall Management Center

Procedure

Step 1 In the navigation bar on the left, click **Administration** > **Migration** > **Migrate FTD to cdFMC**.

Step 2 Click + and choose **On-Prem FMC-managed FTD to** Cloud-Delivered Firewall Management Center.

Note
You can initiate only one migration job at one time.

Step 3 In the **Select On-Prem FMC** area, perform the following:

   a. You can click the **Onboard an FMC** link to onboard the on-premises Firewall Management Center if you have not done already. See Onboard an FMC.
   b. Select the on-premises Firewall Management Center from the available list and click **Next**.

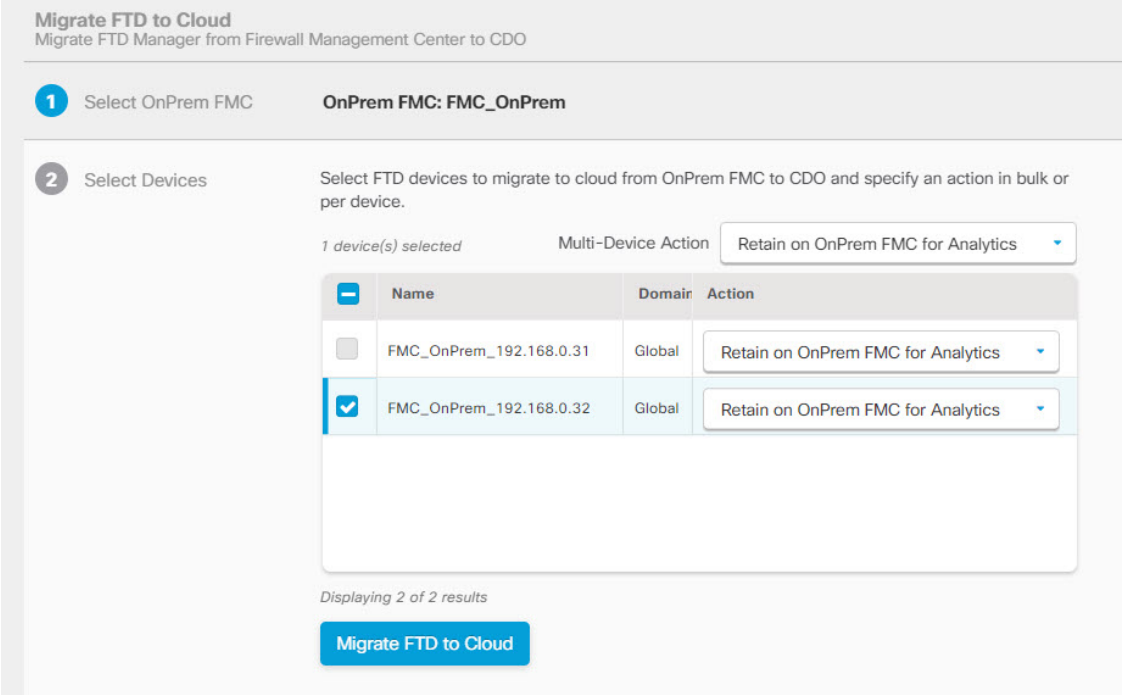
In the **Select Devices** step, you will see the Firewall Threat Defense devices that the selected on-premises Firewall Management Center manages. If a high-availability pair is set up on the on-premise on-premises Firewall Management Center, the high availability node will be shown instead of the active and standby devices.

The **Last Synced time** field indicates the time that is elapsed since the device configuration is synchronized into the on-premises Firewall Management Center. You can click **Sync from OnPrem FMC Now** to fetch the latest device changes.

Step 4 In the **Select Devices** step, perform the following:

a) Select the devices that you want to migrate.

These devices can include standalone, high-availability pairs, or clusters. If you have a high-availability pair or a cluster, select the node that represents it. Also, you can hover your mouse pointer over the cluster to view more details about the control node.



Note

- The devices running on unsupported versions are not available for selection.
- The devices that are registered for analytics only with the on-premises Firewall Management Center or have pending changes to be deployed are not eligible for migration.
- When you select a device that is associated with a site-to-site VPN topology, Security Cloud Control automatically selects its peer devices belonging to either the same topology or a different topology, because all devices in the site-to-site VPN topology must be migrated together for a successful migration to take place. The wizard does not list the extranet devices, if any. However, Security Cloud Control migrates extranet devices.

  The **S2S VPN Topology** column indicates the number of site-to-site VPN topologies in which a selected device participates. You click the topology link to view the topologies and devices that are migrated along with the selected device. This field is not applicable to devices that are not part of the site-to-site VPN topology.
- A high availability pair is presented as a single node. You must select this node to include active and standby devices in the migration.

b) In the **Multi-Device Action** list, you can choose a common action to apply on all devices.
c) In the **Commit Action** column, you can choose one of the following actions for the selected device:

- **Retain on OnPrem FMC for Analytics**: After the migration process is completed, the analytics management for selected Firewall Threat Defense devices is retained on the on-premises Firewall Management Center.
- **Delete FTD from OnPrem FMC**: After the migration process is completed, the selected devices are removed from the on-premises Firewall Management Center and are available for Security Cloud Control to handle the analytics. You must configure the devices to send events to Security Cloud Control for managing analytics. When the devices are deleted from the on-premises Firewall Management Center, they cannot be revoked.

  Important
  For the on-premises Firewall Management Center 1000/2500/4500, when you select devices to migrate, make sure you choose **Delete FTD from OnPrem FMC**. Note that the device is not fully deleted unless you commit the changes or 14 days pass.

Note
The actions that are specified here are committed automatically after the 14 days evaluation period or after the changes are committed manually.

Step 5 (Optional) Check the **Pause migration to review imported shared policies** check box.

When you enable this option, the migration process will pause after the shared access policies like Access Control and NAT policies are staged in the Cloud-Delivered Firewall Management Center. This pause ensures that the evaluation doesn't start and that the device's current state and manager remain unaffected. It gives you ample time to review the imported configuration for accuracy. After you've assessed everything, you can resume the migration during the planned migration interval, which will then kick off the 15-day evaluation period.

Step 6 Check the **Auto deploy to FTDs after successful migration** check box to deploy the migrated configuration automatically to the device after successfully migrating and registering the device with the Cloud-Delivered Firewall Management Center.

However, if you prefer to review and manually deploy the configuration from the Cloud-Delivered Firewall Management Center after successful migration, you can uncheck this option and proceed to the next step.

Step 7 Click **Migrate FTD to cdFMC**.

Step 8 Click **View Migration to Cloud Progress** to see the progress.

What to do next

If you have paused the migration for review, you must manually click **Proceed Migration** to import the remaining configuration. See Proceed Migration Process, on page 82.

You can view the overall and individual status of migration jobs and generate a report when a job is completed successfully. See View a Firewall Threat Defense Migration Job, on page 81.

View a Firewall Threat Defense Migration Job

The migration dashboard provides the status of all migration jobs initiated from the Security Cloud Control. You can expand a specific job to see the status of individual devices associated with that tenant. This helps you keep track of the progress of your migration and identify issues, if any, that need to be addressed.

If you have set up alerts for device workflows, click the notifications icon  to see the alerts that have been triggered during the migration process. Additionally, if you have opted to receive email notifications from Security Cloud Control, you will also receive an email notification regarding alerts, if any.

When a migration job is successful, you have 14 days to test and assess migration changes. If you are convinced about the migration changes, we recommend that you commit the devices manually, and not wait for Security Cloud Control to automatically commit the migration changes. See Commit Migration Changes Manually.

Note that for the on-premises Firewall Management Center 1000/2500/4500, you would have migrated devices from Version 7.4, which is unsupported for general operations. To return the on-premises Firewall Management Center to a supported version you must remove the re-migrated devices, reimage back to Version 7.2.x, restore from backup, and re-register the devices.



Note

- You cannot revoke the actions that are specified in the migration commit window after committing the changes.
- You can cancel the migration during the evaluation period and return the device to the on-premises Firewall Management Center.
- You cannot delete a device from either the on-premises Firewall Management Center or Cloud-Delivered Firewall Management Center during the evaluation period.



Important Changes can be made and deployed to the device using Security Cloud Control during the evaluation period. If you switch device management back to the on-premises Firewall Management Center, Security Cloud Control-specific changes made during the evaluation period is not saved on the device once it is reverted to the source Security Cloud Control tenant. You must deploy the changes from the on-premises Firewall Management Center to the device after reverting the device's manager.

- **Name**: Represents the job name that shows the on-premises Firewall Management Center name and the date and time when the job was initiated.
- **Number of FTDs**: Shows the total number of devices that are being migrated to the cloud.
- **Status**: Shows the status of the job. Expand the job to see the status of individual devices.

  When a job is completed successfully, the **FTD Migration job is successful message** appears in the **Status** column. You can click the tooltip to see the number of days remaining for evaluating the manager.

  You can click Commit Migration Changes to commit the changes manually before the 14 days evaluation period ends.
- **Last Update**: Shows the date and time that are updated only when a change is made to the device.

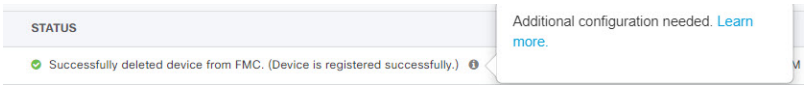
- **Actions**: Click to execute the following actions:
  - **Workflows**: Takes you to the **Workflows** to monitor the job.
  - **Download Report**: Allows you to generate and download a report of every job that is completed successfully. See Generate a Firewall Threat Defense Migration Report, on page 86.
  - **Commit Manager Changes**: Allows you to apply the changes manually to devices before the evaluation period ends. See Commit Migration Changes Manually to Cloud-Delivered Firewall Management Center, on page 83.
  - **Remove Migration Job**: Allows you to remove a completed job. Link is available only for completed jobs. See Delete a Migration Job, on page 86.

After a successful migration, Security Cloud Control deploys the configuration to the device. If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors. If the deployment fails, see the *Best Practices for Deploying Configuration Changes* section of Firepower Management Center Device Configuration Guide X.Y.

Configure Relam Sequence for Identity Policy

If the device contains an identity policy with a Realm or ISE configuration, configure your device as a proxy for Security Cloud Control to communicate with the identity source. The identity policies don't function if Security Cloud Control fails to connect to the Identity Realms.

A tooltip appears in the **Status** column for a device that requires additional configuration.



1. Click the tooltip icon and then click **Learn more**.
2. In the **Configure Proxy** window, click **Configure my realms**.

   To add a proxy sequence, see the *Create a Proxy Sequence* section in the Firepower Management Center Device Configuration Guide, 7.2.

Proceed Migration Process

If the **Pause migration to review imported shared policies** option was checked during the migration setup, the migration process will pause after the shared access policies were staged in the Cloud-Delivered Firewall Management Center. After you've reviewed the staged configurations, you must manually instruct Security Cloud Control to resume the migration, which will then import device-specific configurations such as routing and interfaces, as well as register the Firewall Threat Defense with the Cloud-Delivered Firewall Management Center. Completion of the migration triggers a 14-day evaluation period.

To proceed, go to the migration job page and click **Proceed with migration** for the relevant job.

Remember that deployment from the Cloud-Delivered Firewall Management Center must be done after a successful migration.

Commit Migration Changes Manually to Cloud-Delivered Firewall Management Center

We recommend that you commit migration changes manually if you are convinced with your changes and not waiting for Security Cloud Control to auto commit changes. The **Commit Migration Changes** window shows the remaining days to commit the migration to Cloud-Delivered Firewall Management Center or revert the device to on-premises Firewall Management Center. During the evaluation period, you can modify the actions for selected threat defense devices before committing the changes. Once the changes are committed, you can't revoke the actions.

Note The commit manager changes actions are disabled in the following conditions:

- The 14-day evaluation period has passed.
- The Firewall Threat Defense devices have either been reverted to the on-premises Firewall Management Center or deleted from on-premises Firewall Management Center, in which case, no further actions can be taken.

Procedure

Step 1 In the migration jobs page, click the ... under the **Actions** column of a completed job.

Step 2 Click **Commit Migration Changes**. (This link is available only after a job is completed successfully.)

Step 3 Select a device and in the **Commit Actions** list, choose one of the following actions:

- **Retain on OnPrem FMC for Analytics**: After committing the changes, analytics management for selected Firewall Threat Defense devices is retained on the management center.
- **Delete Firewall Threat Defense from OnPrem FMC**: After committing the changes, the selected devices are removed from the on-premises Firewall Management Center and are available for Security Cloud Control to handle the analytics. You must configure the Firewall Threat Defense to send events to Security Cloud Control for managing analytics. After the Firewall Threat Defense devices are deleted from the on-premises Firewall Management Center, they cannot be revoked.

Note
If you want to revert the device management to on-premises Firewall Management Center, refer to Revert the Firewall Threat Defense Management to On-Premises Firewall Management Center, on page 84.

Step 4 Click **Commit** executes your specified actions immediately without further confirmation.

On the migration jobs screen, you can expand the job to check the progress of the actions specified.

The migrated devices appear on Security Cloud Control's **Security Devices** page. These devices can be managed using the Cloud-Delivered Firewall Management Center portal that is linked to Security Cloud Control. Ensure you deploy the changes to the devices from Cloud-Delivered Firewall Management Center.

Revert the Firewall Threat Defense Management to On-Premises Firewall Management Center

You can revert the device management to the on-premises Firewall Management Center during the evaluation period. This means that the devices will no longer be managed through the Security Cloud Control platform. However, it is important to note that any changes made during the migration process in the Security Cloud Control will not be reflected in the on-premises Firewall Management Center after reverting Firewall Threat Defense management.

Note When the Firewall Threat Defense management has been returned to the on-premises Firewall Management Center , you can begin the migration job again to switch the Firewall Threat Defense management to Security Cloud Control.

Procedure

Step 1 In the migration jobs page, click the ... under the **Actions** column of a completed job.

Step 2 Click **Commit Migration Changes**. (This link is available only after a job is completed successfully.)

Step 3 Select a device and in the **Commit Actions** list, choose **Revert Manager to OnPrem FMC**.

Step 4 Click **Commit** executes your specified actions immediately without further confirmation.

Step 5 Deploy the changes to the device from the on-premises Firewall Management Center.

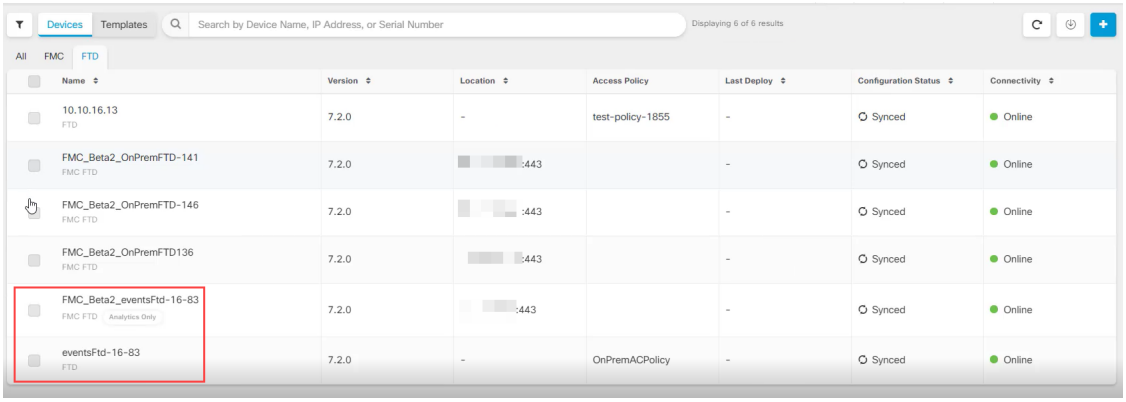
View Migrated Devices

The migrated devices appear on the **Security Devices** page in Security Cloud Control. You can cross-launch and configure the required feature on the Cloud-Delivered Firewall Management Center.

Note The devices on the Cloud-Delivered Firewall Management Center device listing page may show `No-IP` instead of the device's management IP address. Because the device registration uses the NAT ID, the device initiates the process, and therefore, the management IPs aren't discovered or used for the connection. Note that this applies to newly onboarded devices and devices migrated from the on-premises Firewall Management Center.

Analytics Only Firewall Threat Defense Device Example

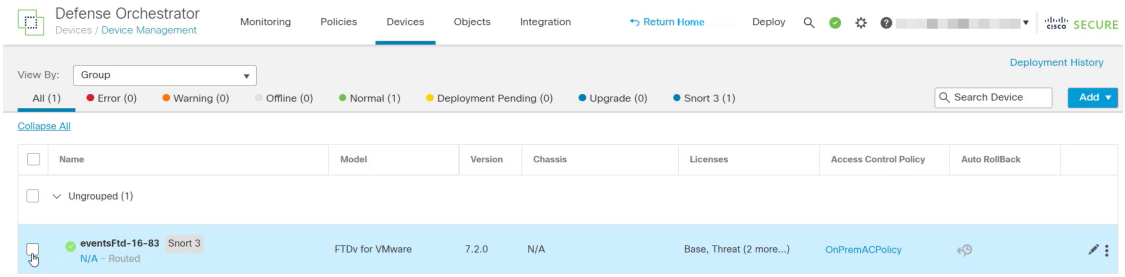
Security Cloud Control creates two instances of the same device that is configured to retain on the Firewall Management Center for analytics.



The device instance with **FMC FTD** and **Analytics Only** labels shows that the Firewall Management Center handles the analytics. The device instance with the **FTD** label indicates that Security Cloud Control manages its configuration.

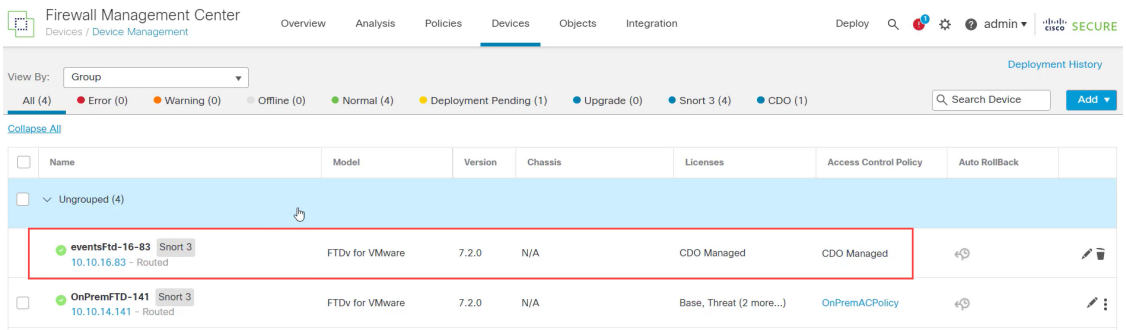
You can manage the configuration of the device using Security Cloud Control. To see the device in the Cloud-Delivered Firewall Management Center, do the following:

Select the device having **FTD** label and in the **Management** pane on the right, click **Device Summary**.



You can view the events from the device in the Firewall Management Center. To see the events, do the following:

1. Select the device having **FMC FTD** and **Analytics Only** labels and on the right, click the **Manage Devices** link.
2. Log on to the on premise Firewall Management Center.
3. Choose **Device** > **Device Management**.



You can't select this device as Security Cloud Control manages the configuration. The Firewall Management Center shows the **Security Cloud Control Managed** label for this device.

To see the live events in the Firewall Management Center, click **Analysis** > **Events**.

Generate a Firewall Threat Defense Migration Report

When a migration job is successful, you can generate and download a report in PDF format to analyze every parameter imported from the on-premises Firewall Management Center to Cloud-Delivered Firewall Management Center. The report provides details of each device associated with the job. Details include information about devices, values of shared policies, objects, routing details, interfaces, network settings, and more.

On the migration jobs page, click the ... under the **Actions** column of a completed job and then click **Download Report.**. You must download a report within a year of the job being triggered.

Delete a Migration Job

If you have completed a migration job and no longer need it to be displayed on the migration page, you can easily remove it by deleting it. This cleans up the migration page and make it easier to navigate.

Attention If you want to delete a migration job during the evaluation period, you must first commit the migration changes or revert the manager to the on-prem management center. Failing to do so may result in an inconsistent state of the on-prem management center, which could be unrecoverable. See Commit Migration Changes Manually.

If you don't have access to your on-premises Firewall Management Center or if it is no longer available and you are blocked from performing a commit or revert, you can delete the job.

Procedure

Step 1 Click **Administration** > **Migration** > **Migrate FTD to cdFMC**.

Step 2 Click the ... under the **Actions** column and then click **Remove Migration Job**.

Step 3 Click **Delete** to confirm your action.

Enable Notification Settings

You can subscribe to get email notifications from Security Cloud Control whenever a device associated with your tenant experiences a specific action when migrating a Firewall Threat Defense device to Security Cloud Control.

Security Cloud Control sends an email if you enable to receive a notification for the following states during migration:

- **Failed**: When a migration job fails.
- **Started**: When a migration job is initiated.
- **Succeeded**: When a migration job is completed successfully.
- **Commit Pending**: When the manager changes have to be committed.

To enable notification settings, see Notification Settings.

Verify Firewall Threat Defense Connectivity with Cloud-Delivered Firewall Management Center

This section provides the commands to determine the Firewall Threat Defense connectivity with the Cloud-Delivered Firewall Management Center.

Check internet connectivity on the device

Execute the **ping system** <*any OpenDNS server address*> command to check whether the device can reach the internet.

1. Connect to the CLI of the device, either from the console port or using SSH.
2. Log in with the Admin username and password.
3. Enter **ping system** <*OpenDNS IPAddress*>.

```
ping system 208.67.222.222
PING 208.67.222.222 (208.67.222.222) 56(84) bytes of data.
64 bytes from 208.67.222.222: icmp_seq=1 ttl=48 time=22.10 ms
64 bytes from 208.67.222.222: icmp_seq=2 ttl=48 time=22.10 ms
64 bytes from 208.67.222.222: icmp_seq=3 ttl=48 time=22.8 ms
64 bytes from 208.67.222.222: icmp_seq=4 ttl=48 time=22.6 ms
^C
--- 208.67.222.222 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 7ms
rtt min/avg/max/mdev = 22.588/22.841/22.995/0.223 ms
```

The above example shows that the device can connect to the internet using the OpenDNS Server IP address. Also, the number of packets transmitted is the same as received, indicating that internet connectivity is available on the device. This shows that the device can reach the internet.

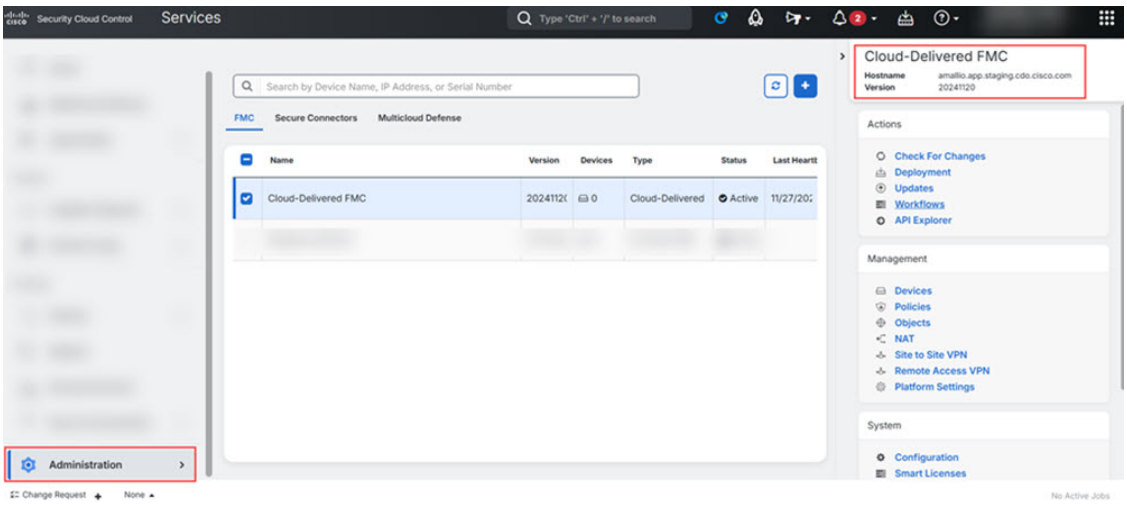
Note If your results don't match, check the internet connection manually.

Check device connectivity with Cloud-Delivered Firewall Management Center

1. Obtain the host name of the Cloud-Delivered Firewall Management Center.
   a. In the Security Cloud Control left pane, click **Administration** > **Integrations** > **Firewall Management Center**.

b. Choose **Cloud-Delivered FMC** to see the Cloud-Delivered Firewall Management Center details on the right pane.

c. In the **Hostname** field, copy only the hostname shown in the following example image.



In the above figure, the highlighted text is the hostname (*Security Cloud Control-acc10.app.us.Security Cloud Control.cisco.com*) of the FMC to be copied.

2. Connect to the CLI of the device, either from the console port or using SSH.

3. Enter **ping system** <*hostname of the FMC*>.

```
ping system cdo-acc10.app.us.cdo.cisco.com
PING cdo-acc10.app.us.cdo.cisco.com (54.187.125.161) 56(84) bytes of data.
^C
--- cdo-acc10.app.us.cdo.cisco.com ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 64ms
```

In the above example, the hostname is resolved with the IP address, indicating your connection is successful. Ignore the "100% packet loss" message shown in the response.

Note If you can't connect to the host, you can rectify the DNS configuration in the CLI using **configure network dns** <*address*>.

Use the Troubleshoot Feature for Failed Migrations

Before you begin

There might be a few instances where, owing to some issues, the migration fails or is not thoroughly successful. To help resolve issues around such failed or partially successful migrations, you can use the inbuilt **Troubleshoot** feature.

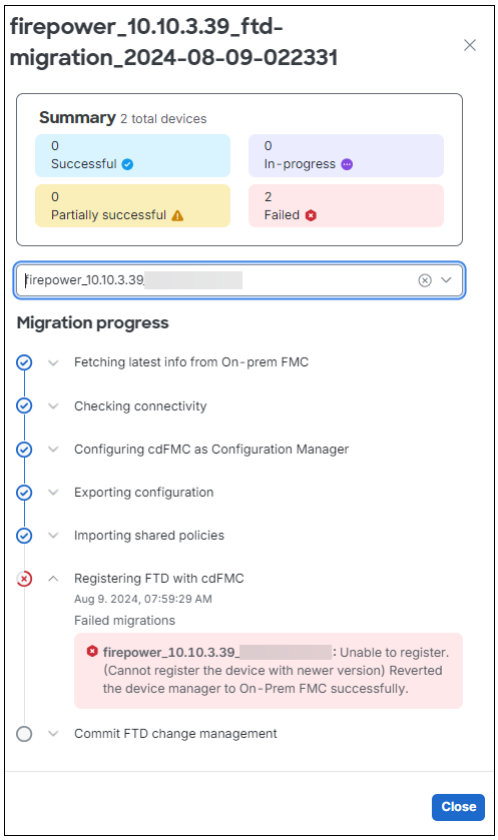
Procedure

Step 1 On the migration jobs page, click the ... under the **Actions** column of a failed or a partially migrated job.

Step 2 Click **Troubleshoot**.

This opens up a sliding window that contains the following information:

- A **Summary** containing additional details about the status of the migration.
- A detailed **Migration progess** listing all the steps in the migration and also indicating the step at which the migration has failed.



Step 3 Based on the stage at which the migration seems to have failed, refer Troubleshoot Firewall Threat Defense Migration to Cloud-Delivered Firewall Management Center for information on how to troubleshoot.

Resolving Firewall Threat Defense Migration to Cloud-Delivered Firewall Management Center Issues

This section details how to troubleshoot specific errors that may occur when migrating Firewall Threat Defense devices to Cloud-Delivered Firewall Management Center.

Migration of Firewall Threat Defense cluster fails

The migration of a Firewall Threat Defense cluster to the Cloud-Delivered Firewall Management Center fails during the migration process. When the failure occurs, Cloud-Delivered Firewall Management Center automatically performs a rollback, returning device management to the original On-Premises Firewall Management Center to maintain configuration integrity and service continuity.

Cause:

The migration failure may occur due to temporary service latency, transient connectivity issues, or other temporary conditions that interrupt the migration process.

Workaround:

- **Retry migration**: Wait for some time and then initiate the migration again from the Secure Firewall Management Center.
- **Contact support**: If the migration continues to fail after multiple attempts, contact Cisco Technical Support and provide the relevant error messages and migration details to assist with troubleshooting.

Configuration export from On-Premises Firewall Management Center failed

Cause:

- Disk space is full or nearly full, preventing the export process from being completed.
- Concurrent operations such as device upgrade, upgrade revert, moving a device between domains, device import and export, device template import and export, or policy analysis are running in on-premises Firewall Management Center.
- Connectivity issues exist between the on-premises Firewall Management Center and Security Cloud Control.

Workaround:

1. Perform these actions according to your requirements:
    - Ensure that at least 10% of disk space is available to accommodate the creation of necessary export files.
    - Ensure that no parallel operations are running during the migration.
    - Ensure that the network connection between on-premises Firewall Management Center and Security Cloud Control is stable and reliable.
2. Retry the migration.
3. If the issue persists, contact Cisco TAC.

Failed to initiate import in Security Cloud Control

Cause:

Import can fail if:

- Disk space is full or nearly full, preventing the import process from being completed.
- Concurrent operations such as device upgrade, upgrade revert, moving a device between domains, device import and export, device template import and export, or policy analysis are running in Cloud-Delivered Firewall Management Center.

Workaround:

1. Ensure that no other operations are running during import and retry migration.
2. If the issue persists, contact Cisco TAC.

Failed to retrieve the most recent device information from On-Premises Firewall Management Center

Cause:

- Network connectivity issues.
- DNS settings.
- Device synchronization issues.

Workaround:

1. Perform these actions according to your requirements:
    - Ensure that the network connection between Security Cloud Control and on-premises Firewall Management Center is stable.
    - Ensure that DNS settings for resolving Firewall Management Center hostnames are correct.
    - Ensure network access for Firewall Threat Defense devices through TCP ports 8305 and 443.
    - Ensure outbound connectivity to port 443 to Security Cloud Control and Cloud-Delivered Firewall Management Center hosts.
    - Ensure that the devices are synchronized without any pending changes.
2. Retry migration.
3. If the issue persists, contact Cisco TAC.

Failed to select Security Cloud Control as the manager for events and analytics

Cause:

Selected on-premises Firewall Management Center for analytics during Firewall Threat Defense migration and committed it or the Security Cloud Control automatically committed it after the 14-day evaluation period.

Workaround:

1. Log in to the on-premises Firewall Management Center and remove the Firewall Threat Defense device from the on-premises Firewall Management Center.

2. Enable the Firewall Threat Defense device to send events to Security Analytics and Logging (SaaS) directly. For more information, see Send Cloud-delivered Firewall Management Center-Managed Event Logs to SAL (SaaS) Using a Direct Connection.
3. Verify if the SSE Connector is running in the Firewall Threat Defense device. To verify, access Firewall Threat DefenseCLI and do the following:

```
# more /ngfw/etc/sf/connector.properties

registration_interval=180

connector_port=8989

connector_fqdn=api-sse.cisco.com
```

4. Verify the connectivity between Firewall Threat Defense and the SSE portal. The following URLs need to be allowed as IPs can change:
   - api.apj.sse.itd.cisco.com (primary SSE cloud API endpoint)
   - est.sco.cisco.com (common endpoint across geographies)
   - mx.apj.sse.itd.cisco.com
   - dex.apj.sse.itd.cisco.com (for customer success services)
   - eventing-ingest.apj.sse.itd.cisco.com (for Cisco XDR and Security Cloud Control services)
   - registration.apj.sse.itd.cisco.com (enables Firewall Threat Defense device registration to the regional Cisco cloud)
5. access Firewall Threat Defense CLI and perform a CURL request to validate tenant details using the following command:

```
root@firepower:/home/admin# curl localhost:8989/v1/contexts/default/tenant
```

6. Verify that the event logs are appearing in Security Cloud Control.

Failed to import configuration from On-Premises Firewall Management Center

Cause:

The On-Premises Firewall Management Center configuration may include policies or object configurations that are unsupported or are conflicting, leading to migration failure.

Workaround:

1. Check the On-Premises Firewall Management Center configuration for unsupported policies or object configurations to ensure compatibility.
2. Retry migration.
3. If the issue persists, contact Cisco TAC.

Failed to complete device discovery

Cause:

Conflicts or unhandled scenarios exist in the device configuration or within Security Cloud Control's internal operations, such as incompatible settings or unsupported conditions that Security Cloud Control does not support.

Workaround:

1. Perform these actions according to your requirements:
    - Inspect the device configuration in On-Premises Firewall Management Center for incompatible settings or unsupported scenarios that Security Cloud Control does not support and make necessary updates.
    - Check system logs for errors and confirm that both the On-Premises Firewall Management Center and the device are functioning properly.
2. Retry migration.
3. If the issue persists, contact Cisco TAC.

Failed to connect the device to Security Cloud Control

Cause:

Possible causes include the device being powered off, network connectivity issues, or a firewall blocking the necessary ports.

Workaround:

1. Perform these actions according to your requirements:
    - Confirm that the device is powered on, and network connection is stable. Look for Firewall issues, blocked ports, and incorrect cabling.
    - Ensure that your device has the correct DNS server configuration for resolving Security Cloud Control hostnames.
2. Retry migration.
3. If the issue persists, contact Cisco TAC.

Failed to register the device, causing the manager to revert to On-Premises Firewall Management Center

Cause:

The device failed to register with Cloud-Delivered Firewall Management Center during migration, causing the system to revert the device manager to On-Premises Firewall Management Center. This may be because of connectivity issues.

Workaround:

1. Perform these actions according to your requirements:
    - Verify DNS settings for resolving Security Cloud Control hostnames.
    - Ensure that network access is enabled for TCP port 8305, allowing the Firewall Threat Defense devices to connect to the Cloud-Delivered Firewall Management Center.

- Allow outbound HTTPS connections on Firewall Threat Defense devices to reach the Cloud-Delivered Firewall Management Center and upload configurations to Security Cloud Control.

2. Retry migration.
3. If the issue persists, contact Cisco TAC.

Failed to unregister the device due to an ongoing deployment

Cause:

Migration failed because a deployment is in progress on the device.

Workaround:

1. Perform these actions according to your requirements:
   - Wait for the ongoing deployment to finish.
   - Confirm that the device is stable, with no pending tasks.
2. Retry migration.
3. If the issue persists, contact Cisco TAC for support.

Report generation failed after successful FTD migration to Cloud-Delivered Firewall Management Center

Cause:

The report generation failure may be because of temporary system glitches.

Workaround:

- Generate the report after some time to overcome the temporary system issues that might have caused the initial failure.
- If the issue persists, contact Cisco TAC.

Failed to initiate FTD migration to Cloud-Delivered Firewall Management Center

Cause:

- Ongoing migration in the source manager.
- Unsupported device version.
- Incorrect DNS and network configuration.

Workaround:

1. Perform these actions according to your requirements:
   - Confirm that no other migration jobs are active in the same source manager.
   - Verify that the devices are on supported versions, have no pending changes to be deployed, and are active.

- Ensure that the DNS settings are correct, and network access is available to communicate with Cloud-Delivered Firewall Management Center.

2. If the issue persists, contact Cisco TAC.

Failed to deploy changes from Security Cloud Control

Cause:

Device migration to Cloud-Delivered Firewall Management Center failed during deployment.

Workaround:

1. Perform these actions according to your requirements:
   - Attempt the deployment again from the **Deployment** page on Cloud-Delivered Firewall Management Center.
   - Check for out-of-band changes, which are changes made directly in the device, and not through Security Cloud Control. If out-of-band changes exist, they may interfere with the deployment. Use Security Cloud Control's **Check for changes** functionality to identify and address out-of-band changes.
   - Ensure that all the configurations are supported and do not conflict with the current device settings.
   - Verify that the device is onboarded and reachable by both the Security Cloud Control and management center, with the necessary internet access and communication capabilities.
2. If the issue persists, contact Cisco TAC.

Failed to commit FTD migration to Cloud-Delivered Firewall Management Center

Cause:

- The 14-day evaluation period has expired.
- Firewall Threat Defense devices have been reverted to or deleted from the On-Premises Firewall Management Center, preventing further actions.

Workaround:

1. Try committing again to address any transient network or device issues.
2. If the issue persists, contact Cisco TAC.

Failed to retain device on On-Premises Firewall Management Center for Analytics

Cause:

If your source manager is an On-Premises Firewall Management Center 1000/2500/4500, it cannot support retaining for Analytics because of the limited resources.

Workaround:

1. Ensure that the **Retain on On-Prem FMC for Analytics** option is selected during the migration process.
2. If the issue persists, contact Cisco TAC.

Failed to revert the device manager to On-Premises Firewall Management Center

Cause:

- The device cannot be reverted if the 14-day evaluation period has expired.
- Firewall Threat Defense devices have been reverted to or deleted from the On-Premises Firewall Management Center, preventing further actions.

Workaround:

1. Perform these actions according to your requirements:
   - Try committing the changes again to address the temporary server or network issues.
   - If changes have occurred in the high availability configuration after migration, reset the devices to their original migration state before attempting to commit.
2. If the issue persists, contact Cisco TAC.

Failed to configure Security Cloud Control as the configuration manager

Cause:

The device already has a configuration manager.

Workaround:

1. If another configuration manager is present, remove it manually and retry the migration.
2. If the issue persists, contact Cisco TAC.

Failed to commit migration; the device is ineligible

Cause:

- The device has pending changes that have to be deployed.
- The device to be migrated is part of FTD high availability (HA), but the current state is invalid, for example, it is in Active/Active state.
- The device to be migrated is connected in Analytics-only mode.
- Cluster migration is supported only for FTDs version 7.4 and above.
- Migration for chassis and MI FTDs on models 3100/4200 is currently not supported.

Workaround:

1. Perform these actions according to your requirements:
   - Verify that the device is up to date with On-Premises Firewall Management Center, then retry migration.
   - Ensure all pending changes on the devices are deployed.
   - Confirm that the Firewall Threat Defense device is in a valid state, such as Active/Standby. If the status isn't correctly reflected in On-Premises Firewall Management Center, use the **Force refresh node status** option on the device listing page to update the correct status on management center.

2. If the issue persists, contact Cisco TAC.

Failed to import site-to-site VPN policies from On-Premises Firewall Management Center

Cause:

The failure may be due to various factors, such as VPN policy configuration or network connectivity issues.

Workaround:

1. Perform these actions according to your requirements:
    - Verify your site-to-site VPN policies in On-Premises Firewall Management Center to ensure there are no configuration errors. If object overrides are used within network objects, manually add these overrides to Security Cloud Control after migration.
    - Ensure that both the On-Premises Firewall Management Center and the targeted device have stable and reliable network connections to Security Cloud Control.
    - If the authentication type is set as **Pre-shared Automatic Key** in On-Premises Firewall Management Center, Security Cloud Control generates a new pre-shared key for the VPN post-migration deployment without disrupting existing tunnels.
2. Retry migration.
3. If the issue persists, contact Cisco TAC.

Failed to migrate as all devices within the topology need to be migrated simultaneously

Cause:

Migrating multiple devices, particularly those within the same network topology, requires synchronizing their migrations to avoid inconsistencies and potential errors. All devices should be migrated within the same timeframe.

Workaround:

1. Perform these actions according to your requirements:
    - Devices registered for Analytics -only with the On-Premises Firewall Management Center or those with pending changes are not eligible for migration.
    - When selecting a device associated with a site-to-site VPN topology, Security Cloud Control automatically selects peer devices from the same or different topology, as all devices in the topology must be migrated together for success. Extranet devices, if any, are not listed by the wizard.
    - The **S2S VPN Topology** column shows the number of site-to-site VPN topologies a device is part of. Click the topology link to view the topologies and devices migrating with the selected device. This field does not apply to devices outside the site-to-site VPN topology.
    - A high-availability pair is shown as a single node. Select this node to include both active and standby devices in the migration.
2. If the issue persists, contact Cisco TAC.

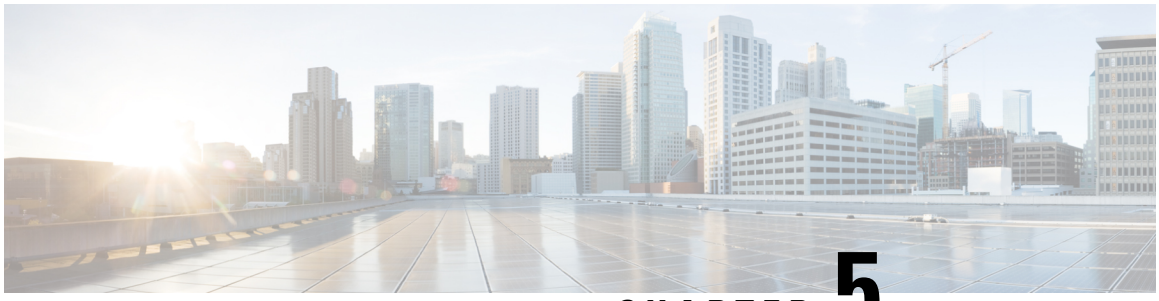
HTTP status code 201 (created) found in On-Premises Firewall Management Center response

Cause:

The Secure Device Connector (SDC) version is not compatible.

Workaround:

1. Ensure that the SDC is upgraded to version **202205191350** or later.
    - In the left pane, click **Administration** > **Integrations** > **Secure ConnectorsAdministration** > **Secure Connectors**.
    - Click the SDC to view the existing version in the **Details** pane on the right.
    - Update your Secure Device Connector.



CHAPTER 5

Device Registration

You canmanage devices in the Cloud-Delivered Firewall Management Center.

- Log into the Command-Line Interface on the device, on page 99
- Device registration management, on page 101

Log into the Command-Line Interface on the device

You can log directly into the command-line interface on Firewall Threat Defense devices. If this is your first time logging in, complete the initial setup process using the default **admin** user; see Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI, on page 51 .

For zero-touch provisioning , if you must access the Firewall Threat Defense CLI and run through the setup script, answer **n** when prompted: `Do you want to configure IPv4? (y/n) [y]:` and `Do you want to configure IPv6? (y/n) [y]:` . You also must accept the default local manager: `Manage the device locally? (yes/no) [yes]:` . These settings will preserve zero-touch provisioning capability.

For the Secure Firewall 200 , the device only supports up to three concurrent CLI sessions. For example, you can have one console session and two SSH sessions to the Management interface (this limitation is separate from SSH to a data interface). If you already have three active SSH sessions and then connect to the console, the console connection is allowed because console access will never be blocked.

Note If a user makes three consecutive failed attempts to log in to the CLI via SSH, the system terminates the SSH connection.

Before you begin

- Create additional user accounts that can log in to the CLI using the **configure user add** command.
- If you get unreadable characters when connecting to the console port, verify the port settings. If they are correct, try the cable with another device using the same settings. If the cable is good, you might need to replace the hardware for the console port. Also consider trying a different workstation to make the connection.

Procedure

Step 1 Connect to the Firewall Threat Defense CLI, either from the console port or using SSH.

You can SSH to the management interface of the Firewall Threat Defense device. You can also connect to the address on a data interface if you open the interface for SSH connections. SSH access to data interfaces is disabled by default. See Configure Secure Shell (SSH) access, on page 1262 to allow SSH connections to specific data interfaces.

For physical devices, you can directly connect to the console port on the device. See the hardware guide for your device for more information about the console cable. Use the following serial settings:

- 9600 baud
- 8 data bits
- No parity
- 1 stop bit

The CLI on the console port is FXOS (with the exception of the ISA 3000, where it is the regular Firewall Threat Defense CLI) . Use the Firewall Threat Defense CLI for basic configuration, monitoring, and normal system troubleshooting. See the FXOS documentation for information on FXOS commands.

Step 2 Log in with the **admin** username and password.

Example:

```
firepower login: admin
Password:
Last login: Thu May 16 14:01:03 UTC 2019 on ttyS0
Successful login attempts for user 'admin' : 1
firepower#
```

Step 3 If you used the console port, access the Firewall Threat Defense CLI.

connect ftd

Note
This step does not apply to the ISA 3000.

Example:

```
firepower# connect ftd
>
```

Step 4 At the CLI prompt (>), use any of the commands allowed by your level of command line access.

To return to FXOS on the console port, enter **exit** .

Step 5 (Optional) If you used SSH, you can connect to FXOS.

connect fxos

To return to the Firewall Threat Defense CLI, enter **exit** .

Step 6 (Optional) Access the diagnostic CLI:

system support diagnostic-cli

Use this CLI for advanced troubleshooting. This CLI includes additional **show** and other commands.

This CLI has submodes: user EXEC mode, privileged EXEC mode , and recovery-config mode . More commands are available in privileged EXEC mode than user EXEC mode. To enter privileged EXEC mode, enter the **enable** command; press enter without entering a password when prompted.

Example:

```
> system support diagnostic-cli
firepower> enable
Password:
firepower#
```

To use recovery-config mode, see Access recovery-config mode in the diagnostic CLI, on page 1146 .

To return to the regular CLI, type **Ctrl-a** , **d** .

Device registration management

Register and unregister devices to the Cloud-Delivered Firewall Management Center.

About the device management page

The **Devices** > **Device Management** page provides you with range of information and options.

You can enable or disable the **New device management UI** toggle button to switch between the legacy and enhanced view in the **Device Management** page.

Figure 8: Toggle button to switch between legacy and enhanced view



- **View By**—View devices based on group, licenses, model, version, or access control policy.
- **Device State**—View devices based on state (**Error**, **Warning**, etc.). You can click on a state icon to view the devices belonging to it. The number of devices belonging to the states are provided within brackets.
- **Search Device**—Search for a device by device name, host name, or IP address.
- **Add**—Add devices and other manageable components.
- Columns—Click the column head to sort by that column.
  - **Name**
  - **Model**
  - **Version**
  - **Chassis**—For supported models, click **Manage** to bring up the integrated Chassis Manager. For the Firepower 4100/9300, the link cross-launches the Firewall Chassis Manager.

- **Licenses**
- **Access Control Policy**—Click on the link in the Access Control Policy column to view the policy that is deployed to the device.
- **Auto-Rollback**—Shows whether auto-rollback of the configuration is enabled ( ) or disabled ( ) if the deployment causes the management connection to go down. See Edit deployment settings, on page 1192.

- **Edit**—For each device, use the **Edit** ( ) icon to edit the device settings.

  You can also just click on the device name or IP address.

- **More**—For each device, click the **More** ( ) icon to execute other actions:
  - **Packet Tracer**—To navigate to the packet tracer page for examining policy configuration on the device by injecting a model packet into the system.
  - **Packet Capture**—To navigate to the packet capture page, where, you can view the verdicts and actions the system takes while processing a packet.
  - **Revert Upgrade**—To revert the upgrade and configuration changes that were made after the last upgrade. This action results in restoring the device to the version that was before the upgrade.
  - **Health Monitor**—To navigate to the device's health monitoring page.
  - **Troubleshoot Files**—Generate troubleshooting files, where you can choose the type of data to be included in the report.
  - **Generate Template from Device**—Generate a new device template from a registered device. The new template has the same configuration as the device from which it is generated. You can generate a new device template from standalone and HA devices. However, if you generate a template from HA devices, the new template will not contain the failover configurations.

Add a Device Group

The Cloud-Delivered Firewall Management Center allows you to group devices so you can easily deploy policies and install updates on multiple devices. You can expand and collapse the list of devices in the group.

If you add the primary device in a high-availability pair to a group, both devices are added to the group. If you break the high-availability pair, both devices remain in that group.

Procedure

Step 1 Choose **Devices** > **Device Management** .

Step 2 From the **Add** drop-down menu, choose **Add Group** .

To edit an existing group, click **Edit** ( ) for the group you want to edit.

Step 3 Enter a **Name** .

Step 4 Under **Available Devices** , choose one or more devices to add to the device group. Use Ctrl or Shift while clicking to choose multiple devices.

Step 5 Click **Add** to include the devices you chose in the device group.

Step 6 Optionally, to remove a device from the device group, click **Delete** ( ) next to the device you want to remove.

Step 7 Click **OK** to add the device group.

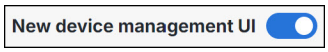
Device management page enhancements

Device management page is enhanced to a new user interface that can be toggled with the legacy version. User interface improvements offer enhanced features including improved search functionality, device status monitoring, and centralized troubleshooting tools.

Device management page enhancement features

Navigation: **Devices** > **Device Management**

Enable the **New device management UI** toggle button to switch between the legacy and the enhanced **Device Management** page.



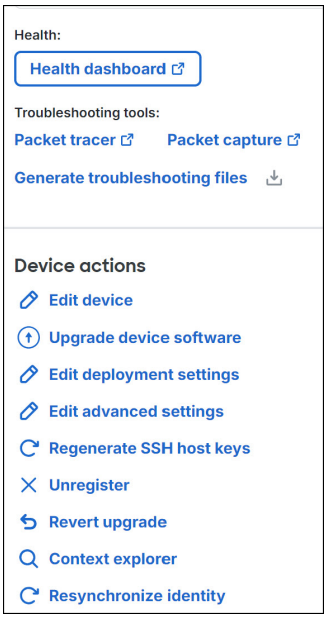
Note The enhanced user interface is in a preliminary stage and does not include all the legacy user interface features. You can switch back to the legacy **Device Management** user interface by clicking the same toggle button to access the features that are not yet supported in the new user interface.

Key enhancements in the new user interface:

- **Improved search functionality**: Provides a more powerful and granular search experience. Unlike the legacy user interface, which was limited to searching by device name, host name, or IP address, you can now use multiple device-related criteria, to achieve highly refined results.
- **Device status banner**: A status banner at the top of the **Device Management** page displays the number of devices in various states (**Normal**, **Error**, **Offline**). A color-coded legend dynamically reflects these statuses for at-a-glance monitoring.
- **Performance improvements**: The enhanced user interface page loads faster than the legacy version. Pagination has been introduced at the bottom of the page, with each page supporting up to 1,000 devices by default.
- **Device actions and bulk operations**: Details of selected devices are displayed on a side panel in a more user-friendly format. Multiple devices can be selected simultaneously enabling device bulk actions from a centralized location.
- **Centralized troubleshooting panel**: Troubleshooting information has been consolidated into a single pane under Troubleshooting tools, making it easier to access diagnostic tools such as **Packet Tracer** and **Packet Capture**.

Device settings

In the new device management page view, select the device to open the side pane. Click the option you want:



- Under **Health and troubleshooting**, click **Health dashboard** to navigate to the device's health monitoring page.
- Under **Troubleshooting tools**, click:
  - **Packet Tracer**—To navigate to the packet tracer page for examining policy configuration on the device by injecting a model packet into the system.
  - **Packet Capture**—To navigate to the packet capture page, where, you can view the verdicts and actions the system takes while processing a packet.
  - **Generate troubleshooting files**—To generate diagnostic files for the device. You can choose the type of data to be included in the report.
- Under **Device actions**, click:
  - **Edit**—For each device, use the **Edit** (✎) icon to edit the device settings.
  - **Upgrade device software**—Install a software update on this device.
  - **Edit deployment settings**—Configure automatic rollback if a deployment disrupts management connectivity.
  - **Edit advanced settings**—Modify advanced device configuration settings.
  - **Regenerate SSH host keys**—Generate new SSH host keys for this device.
  - **Unregister**—Remove this device from the management center.
  - **Revert upgrade**—Restore the device to the software version and configuration from before the last upgrade.
  - **Content explorer**—Open Context Explorer to investigate device data.
  - **Resynchronize identity**—Synchronize identity data for this device.

Use the enhanced filters in device management page

This task enables you to quickly find specific devices using various search criteria and save frequently used searches for future use.

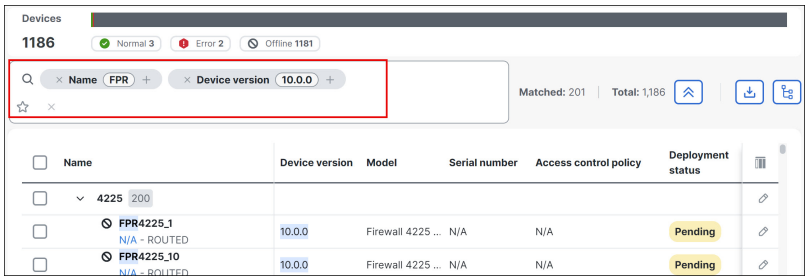
You can use different criteria to instantly search for the devices you want.

Procedure

Step 1 In the **Device Management** page, click the **Type to search** field to perform one of the following actions:

- Enter the filter criteria to search the device. Suggestions are displayed as you type.
- Select a filter and then enter a value in the field. You can add multiple criteria to refine your search results further.

In the example shown in the figure, the filter searches for devices with names that contain "FPR" andare of version of "10.0.0".



Step 2 Click the **Favourite Searches** (☆) icon and do one of the following:

- To save a new search, specify a search name and click **Save as new**.
- To overwrite a saved search, click the **Edit** () icon on next to the saved search that you want to overwrite, and click **Overwrite**.

Note

You can also use the **Favourite Searches** (☆) icon to quickly retrieve the previously executed searches.

The enhanced filters display devices that match your search criteria, and saved searches are available for future use through the Favourite Searches feature.

Add a device group

Create device groups to organize and manage multiple devices efficiently, enabling streamlined policy deployment and update installation across grouped devices.

The Management Center allows you to group devices to easily deploy policies and install updates on multiple devices. You can expand and collapse the list of devices in the group. The new device management user

interface provides enhanced functionalities in the context of device groups, improving the way you organize, manage, and interact with them.

The **Manage device groups** dialog box provides two main sections:

- Ungrouped devices: This section lists all the devices that are currently not assigned to any group. A search bar, labeled **Search devices to move** enables you to quickly locate specific devices by device name or device IP address. The count that is displayed represents the number of ungrouped devices, and not the total number of devices managed by the system.
- Groups: This section displays all the existing device groups. A search bar, **Search by group or device name**, enables you to find specific groups or devices within groups.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Enable the **New Device Management UI** toggle button and click the **Manage Groups** () icon.

Step 3 Click **Create Group**.

Step 4 Enter a name and click .

The group is created.

Step 5 Select the devices that you want to add to the newly created group.

Step 6 Under **Groups** locate the newly created group.

You can use the **Search by group or device name** field to search for the group.

Step 7 Click **Add selection to group**.

The selected devices are added to the group.

Note
You can also delete a group. However, deleting a group does not remove the devices from the Management Center. Instead, all the devices assigned to the deleted group will be automatically moved back to the Ungrouped devices section.

A new device group is created with the selected devices assigned to it. The devices appear under the group in the Groups section and are removed from the Ungrouped devices section.

Move devices to another group

Moving devices to another group allows you to reorganize your device management structure and ensure devices are properly categorized for administrative purposes.

The device must first be moved to Ungrouped Devices. Next, move it to the group you want.

Procedure

Step 1 Click the **Manage Groups** icon.

Step 2 Under **Groups**, expand the group from which you want to move the device.

Step 3 Click the X button to remove the device from the group.

The device is moved to **Ungrouped Groups**.

Step 4 Select the device that you want to add.

Step 5 Under **Groups** locate the group.

You can use the **Search by group or device name** field to search for the group.

Step 6 Click **Add selection to group**.

The selected device is added to the group.

Managed device list download

A managed device list download is a data export feature that enables administrators to obtain device inventory information in multiple file formats for offline analysis and reporting purposes.

Download process

Click  to download the device list in CSV or PDF format. Click **Download CSV** or **Download PDF** to download the report in the corresponding format.

Shut down or restart the device

It's important that you shut down your system properly. Simply unplugging the power or pressing the power switch can cause serious file system damage. Remember that there are many processes running in the background all the time, and unplugging or shutting off the power does not allow the graceful shutdown of your firewall.

See the following task to shut down or restart your system properly.

Note After restarting your device, you may see an error that the management connection could not be reestablished. In some cases, the connection is attempted before the Management interface on the device is ready. The connection will be retried automatically and should come up within 15 minutes.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device that you want to restart, click **Edit** ( ) .

Step 3 Click **Device**.

Step 4 To restart the device:

a) Under **System**, click **Restart Device** ( ) .
b) When prompted, confirm that you want to restart the device.

Step 5 To shut down the device:

a) Under **System**, click **Shut Down Device** ( ) in the **System** section.
b) When prompted, confirm that you want to shut down the device.
c) If you have a console connection to the firewall, monitor the system prompts as the firewall shuts down. You will see the following prompt:

```
System is stopped.
It is safe to power off now.
Do you want to reboot instead? [y/N]
```

If you do not have a console connection, wait approximately 3 minutes to ensure the system has shut down.

For the ISA 3000, when shutdown is complete, the System LED will turn off. Wait at least 10 seconds before you remove the power.

Download the managed device list

You can download a report of all the managed devices.

Before you begin

To perform the following task, you must be an Admin user.

Procedure

Step 1 Choose **Devices** > **Device Management** .

Step 2 Click the **Download Device List Report** link.

Step 3 You can download the device list in CSV or PDF format. Choose **Download CSV** or **Download PDF** to download the report.

Migrate Firewall Threat Defense devices

The Secure Firewall Threat Defense model migration wizard enables you to migrate configurations from an earlier Firewall Threat Defense model. After the migration, all routing and interface configurations from the source Firewall Threat Defense device are available in the target Firewall Threat Defense.

Supported devices and capabilities

The wizard supports multiple models as source and target devices, for more information see Supported Devices for Migration, on page 109.

When you migrate Firepower 4100 and 9300 Series devices to the supported models, you can now configure interface attributes according to your requirements. You can map the source device interfaces to the target device interfaces. The migration locks the source and target devices.

Supported Devices for Migration

Documents the supported migration paths and device compatibility for Cisco Firepower and Secure Firewall models.

Supported Migration Paths

The following table lists the supported target Firewall Threat Defense models that you can migrate to from your source Firewall Threat Defense model.

Source Devices	Source Device Version	Target Devices	Target Device Version
Cisco Firepower 1010 Series: 1010, 1010E	7.3.x and later	Cisco Secure Firewall 1200 Series:1210CE, 1210CP, 1210CX	7.6 and later
Cisco Firepower 1010 Series: 1010, 1010E	7.3.x and later	Cisco Firewall 200 Series: 220	10.0 and later
Cisco Firepower 1100 Series: 1120, 1140, 1150	7.3.x and later	Cisco Secure Firewall 3100 Series: 3105, 3110, 3120, 3130, 3140	7.4.1 and later
Cisco Firepower 2100 Series: 2110, 2120, 2130, 2140	7.3.x and later	Cisco Secure Firewall 3100 Series: 3105, 3110, 3120, 3130, 3140	7.4.1 and later
Cisco Firepower 4100 Series: 4112, 4115, 4125, 4145	7.3.x and later	Cisco Secure Firewall 3100 Series: 3105, 3110, 3120, 3130, 3140	7.4.1 and later
		Cisco Secure Firewall 4200 Series: 4215, 4225, 4245	7.4.1 and later
		Cisco Secure Firewall 6100 Series: 6160, 6170	10.0 and later

Source Devices	Source Device Version	Target Devices	Target Device Version
Cisco Firepower 9300 Series: SM-40, SM-48, SM-56	7.3.x and later	Cisco Secure Firewall 3100 Series: 3105, 3110, 3120, 3130, 3140	7.4.1 and later
		Cisco Secure Firewall 4200 Series: 4215, 4225, 4245	7.4.1 and later
		Cisco Secure Firewall 6100 Series: 6160, 6170	10.0 and later
Cisco ASA 5500 Series: 5508, 5516	7.0.x	Cisco Secure Firewall 1200 Series:1210CE, 1210CP, 1210CX	7.6 and later
Cisco Firepower 1100 Series: 1120, 1140, 1150	7.3.x and later	Cisco Secure Firewall 1200 Series:1210CE, 1210CP, 1210CX	7.6 and later
Cisco Firepower 1100 Series: 1120, 1140, 1150	7.3.x and later	Cisco Secure Firewall 1200 Series: 1230, 1240, 1250	7.7.x and later
Cisco Firepower 2100 Series: 2110, 2120, 2130, 2140	7.4.x	Cisco Secure Firewall 1200 Series:1210CE, 1210CP, 1210CX	7.6 and later
Cisco Firepower 2100 Series: 2110, 2120, 2130, 2140	7.4.x	Cisco Secure Firewall 1200 Series: 1230, 1240, 1250	7.7.x and later

License for Firewall Threat Defense migration

You must meet licensing requirements for Firewall Threat Defense migration.

- Your Smart Licensing account must have the license entitlements for the target device.
- You must register and enroll the device with the Smart Licensing account. The migration copies the source device licenses to the target device.

Prerequisites for migration

This reference provides comprehensive prerequisites that must be satisfied to successfully migrate configurations from a source device to a target device in Secure Firewall Management Center environments.

General device prerequisites

- Register the source and the target devices to the Cloud-Delivered Firewall Management Center.
- Ensure that the target device is a newly registered device without any configurations.
- Source and target devices must be in the same state and modes:

- Firewall mode: Routed or Transparent
- Compliance mode (CC or UCAPL)
- Management state

  Devices must have the same type of manager access interfaces (management interface or data interface).
- Multi-instance mode or appliance mode

- Ensure that you have permission for modifications on the devices.
- Ensure that the configurations on the source device are valid and have no errors.
- Deployment, import, or export tasks must not run on either of the devices during the migration. The source device can have pending deployments.

Prerequisites for change management

- Ensure that source and target devices are not locked by a change management ticket.
- Ensure that shared policies assigned to the source device are not locked by a change management ticket.

Prerequisites for devices in multi-instance mode

- Ensure that the source and target devices are in multi-instance mode.
- Manually migrate the chassis configurations. Create instances before migrating the instance configuration to the target instances. The target device must have compatible interfaces. For example, on the target device, you must create EtherChannel interfaces, and also create tagged, untagged, dedicated, or shared interfaces for these interfaces on the target device.

Prerequisite for devices with out-of-band configurations

- Ensure that you acknowledge out-of-band changes and match the configurations within the Cloud-Delivered Firewall Management Center. You cannot migrate devices with these configurations. To view out-of-band configurations:

  1. Choose **Devices** > **Device Management**.
  2. Click the edit icon next to the device and click the **Interfaces** tab.

  or

  1. Click the edit icon next to the device and click the **Devices** tab.
  2. Verify the **Out-of-band configuration status** in the **Health** tile.

Prerequisites for devices with manager access interfaces

Ensure that the devices are not in Data Transit or Management Transit states. You cannot migrate if devices are in these states.

- Data Transit state: Device state when the manager access interface changes from data interface to management interface without deploying the changes on the device.
- Management Transit state: Device state when the manager access interface changes from management interface to data interface without deploying the changes on the device.

Prerequisite for devices with merged management and diagnostic interfaces

Ensure that the target device is always in merged mode.

Configurations that the wizard migrates

The migration wizard copies configurations from the source device to the target device, including licenses, interfaces, policies, routing settings, and other essential device configurations.

General configurations

The migration wizard copies these configurations from the source device to the target device:

- Licenses
- Interface configurations
- Inline sets configurations
- Routing configurations
- DHCP and DDNS configurations
- Policies
- Associated objects and object overrides
- Platform settings
- Remote branch deployment configurations

Policy configurations

The migration wizard copies these policy configurations from the source device to the target device:

- Health policy
- NAT and QoS policy
- Remote access VPN policy
- FlexConfig policy
- Access control and prefilter policy
- IPS and DNS policy
- Malware and File policy
- SSL and Identity policy
- Shared policy

Routing configurations

The migration wizard copies these routing configurations from the source device to the target device:

- ECMP
- BFD
- OSPFv2/v3
- EIGRP
- RIP
- BGP
- Policy Based Routing
- Static and Multicast Route
- Virtual Router

Interface configurations

The migration wizard copies these interfaces from the source device to the target device:

- Physical interfaces
- Sub-interfaces
- EtherChannel interfaces
  - On a standalone device, the wizard copies the EtherChannels from the source device to the target device.
  - For devices in multi-instance mode, you must create EtherChannels on the chassis and assign them to the instance.
- Bridge group interfaces
- VTI interfaces
- VNI interfaces
- Loopback interfaces
- VXLAN tunnel endpoint (VTEP) interfaces

The migration wizard retains the device group of the target device.

Guidelines and limitations for migration

Follow these guidelines to ensure successful migration of secure firewall threat defense models:

- **For devices in multi-instance mode**:

  During migration, ensure that you map the interfaces according to the table below:

Source Device	Target Device
Physical interface	Physical interface
EtherChannel interface	EtherChannel interface
Supervisor-provisioned subinterface	Supervisor-provisioned subinterface
Tagged interface	Tagged interface
Untagged interface	Untagged interface
Shared interface	Shared and dedicated interface
Dedicated interface	Dedicated interface

You cannot map a supervisor-provisioned subinterface to a subinterface created by an instance.

- **For HA devices**, you can migrate:
  - Source HA device to target HA device.
  - Source HA device to target standalone device.
- **For devices in remote branch deployment**:
  - Map the source manager access interface to the target manager access interface.
  - Ensure that the manager access interfaces of the source and target Cloud-Delivered Firewall Management Centers are of the same IP address type (static or DHCP).
  - Both manager access interfaces must have IPv4 or IPv6 addresses.
  - If the manager access interfaces have static IP addresses, ensure that they are in the same subnet.
- **For Snort**:

  By default, after migration, the target device will use Snort 3, even if the source device uses Snort 2.
- **For devices using diagnostic interfaces**:

  Only merged management interfaces are available on the target devices after migration.

Limitations

Be aware of these limitations when using the migration wizard:

- The migration wizard does not migrate:
  - Site-to-site VPN policies
  - SNMP device configurations for Firepower 2100 Series

    After the migration, you can configure SNMP using the platform settings for the device.
- You can perform only one migration at a time.
- Remote access VPN trustpoint certificates are not enrolled after migration.

- For HA devices:
  - Target device: You cannot migrate a standalone device to an HA device.
- Clustering is not supported.
- For devices in remote branch deployment:
  - The wizard does not migrate a single WAN manager access data interface to a dual WAN manager access data interface.

Migrate a secure Firewall Threat Defense

Migrate a Secure Firewall Threat Defense device from a source to a target Firewall, preserving configurations and minimizing disruptions.

Before you begin

Ensure you review Prerequisites for migration, on page 110 and Guidelines and limitations for migration, on page 113 .

Procedure

Step 1 Choose **Devices** > **Device Management** .

Step 2 Click **Migrate** in the top right corner of the page.

Step 3 In **Select source and target devices** :

a) From the **Source device** drop-down list, choose a device.

b) From the **Target device** drop-down list, choose a device.

The source and target devices can have these tags:

- Routed: Devices in routed firewall mode.
- Transparent: Devices in transparent firewall mode.
- Container: Devices in multi-instance mode.
- High Availability: Devices in high availability mode.
- Analytics Only: Devices managed by Security Cloud Control and the Cloud-Delivered Firewall Management Center only recieves and displays the events (analytics-only Cloud-Delivered Firewall Management Center).

If the device is part of an HA pair, only the HA pair name appears.

Step 4 Click **Next** .

Step 5 (Only for Firepower 4100 and 9300 Series devices in appliance mode) In **Chassis manager details** :

a) Check the **Skip chassis manager** check box, if required.

b) In the **Chassis hostname or IP address** field, enter the values.

Note

- Verify that the Secure Firewall Chassis Manager is reachable from the Cloud-Delivered Firewall Management Center .

- Ensure you select the correct chassis manager for the source device, as Cloud-Delivered Firewall Management Center does not validate your choice.

c) Click **Verify certificate** to verify the chassis manager's certificate.
d) In the **Username** and **Password** fields, enter the credentials of the chassis manager.

Step 6 Click **Next** .

Step 7 In **Configure interfaces** :

By default, the source and target interfaces are mapped using the interface hardware name. You must map named interfaces, logical interfaces, and interfaces that are part of other interfaces. Mapping of all other interfaces is not mandatory. The wizard creates the logical interfaces according to the interface mapping that you provide.

You cannot map interfaces that are part of an HA failover configuration. These interfaces are disabled in the wizard.

Only data interfaces are available for interface mapping. Management, eventing, and diagnostic interfaces are not available for the interface mapping.

Firepower 4100 and 9300 Series devices in appliance mode :

For these devices, the Cloud-Delivered Firewall Management Center fetches interface attributes such as speed, duplex, and auto-negotiation from the chassis manager.

a) Click one of the following options to configure these interface attributes on the target device:

- **Retain target device values** : (Default) Retains the interface attributes configured on the target device.
- **Copy from source device** : Copies the interface attributes from the source device.

  This option is enabled only when Cloud-Delivered Firewall Management Center successfully connects to the chassis manager. We recommend that you use this option. The speed, duplex, and auto-negotiation values of physical interfaces are set to default values if they are incompatible in the target device.
- **Customize device values** —Allows you to configure the values of the required interface attributes on the target device.

b) To change the interface mapping from the default ones, choose an interface from the **Mapped interface** drop-down list.
c) For EtherChannels, you can configure interface attributes and click **Add member interface** to add member interfaces.

   Interface attributes of an EtherChannel is configured based on the first member interface's interface attributes.You can add up to 16 member interfaces.

Firepower 1100 and 2100 Series devices, and Firepower 4100 and 9300 Series devices in multi-instance mode :

For these devices, you must map the source device interfaces to target device interfaces.

For Firepower 4100 and 9300 Series devices in multi-instance mode, you can only perform the interface mapping and you cannot configure the interface attributes such as speed, duplex, auto-negotiation, and FEC mode.

If you want to change the interface mapping from the default ones, choose an interface from the **Mapped interface** drop-down list.

Click **Reset** to configure the default interface mappings. For example, the wizard maps Ethernet1/1 in the source device to Ethernet1/1 in the target device.

The interfaces can have the following tags:

- Tagged: Physical interfaces on the chassis.
- Untagged: Physical interfaces on the chassis that have sub-interfaces.
- Dedicated: Interfaces that are assigned to specific instances and are not shared across multiple instances.
- Shared: Interfaces that are shared by multiple instances.
- Manager access: Data interface is the manager access interface.

Check the **Ignore warning** check box, if required.

Step 8 Click **Next** .

Step 9 Click **Submit** to start the migration.

To view the migration status, click **Notifications** (Message Center), and then click the **Tasks** tab.

A **Device Model Migration** report is generated after the migration is completed. You will see a link to this report in the **Notifications** > **Tasks** page.

What to do next

After a successful migration, you must complete these tasks:

- Review the recommendations in Best practices for Threat Defense device migration, on page 117 .
- Validate the configurations.
- Deploy the configurations on the device.

In case of a migration failure, the target device is rolled back to the initial state.

Best practices for Threat Defense device migration

After a successful migration, perform these actions before deployment to ensure proper device configuration and functionality.

- IP addresses of the interfaces are copied to the target device from the source device. Change the IP addresses of the target device interfaces if the source device is live
- Ensure that you update your NAT policies with the modified IP addresses.
- Configure the interface speeds if they are set to default values after migration.
- Re-enroll the device certificates, if any, on the target device.
- (Optional) Configure SNMP for Firepower 1100 and 2100 using the platform settings for the device.
- (Optional) Configure remote branch deployments.

If the source or target device had manager access through a data interface, after the migration, the manager access will be lost. Update the manager access configuration on the target device. For more information, see the *Change the Manager Access Interface from Management to Data* topic in the Cisco Secure Firewall Management Center Device Configuration Guide or the Online Help.

- Configure site-to-site VPN, if required. These configurations are not migrated from the source device.
- View the deployment preview before the deployment. From the **Deploy** drop-down menu, click **Advanced Deploy**, and then click the **Preview** ( ) icon for the device.
- Monitor the health of the device in the health monitor (choose **System** ( ) > **Health** > **Monitor**). After migration, the health policy of the source device becomes the health policy of the target device. You can also configure a new health policy for the device.

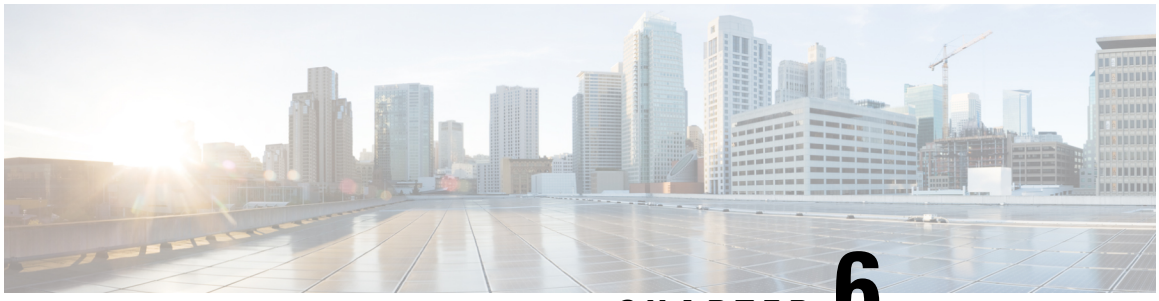
  After migration, the device monitoring dashboard may temporarily display redundant colored lines because the device has different UUIDs before and after migration. This redundancy appears only during the migration time. An hour after migration, the dashboard will show a single line per metric.

Use the Security Cloud Control command line interface tool

You can use the command line interface (CLI) to troubleshoot Firewall Threat Defense and other supported device types within Security Cloud Control.

Additional information

For more information, see the Security Cloud Control configuration guide.



CHAPTER 6

Device Registration Using Device Templates

You can add devices to the Cloud-Delivered Firewall Management Center using device templates.

- Device registration using device templates, on page 119
- Prerequisites for device registration using device templates, on page 121
- Licenses for device registration using device templates, on page 122
- Guidelines for device registration using device templates, on page 123
- Configure device templates, on page 126
- Use templates with devices, on page 145
- Monitor device templates, on page 147
- Troubleshoot device templates, on page 149

Device registration using device templates

You can use device templates to apply configuration changes to multiple devices with different interface configurations, and clone configuration parameters from existing devices.

Configuration capabilities

Device templates let you pre-configure settings and apply them at the time of registration. When you register a device using basic initial configuration, you can apply limited configurations such as the access control policy and licenses. You must then configure other device settings such as interfaces, routing, and site-to-site VPN configurations individually after device registration. Values that need to be unique per device, such as IP addresses, can be defined using variables and network object overrides that you define at registration.

You can also configure site-to-site VPN connections in a device template. These configurations define the site-to-site VPN topologies that a device should be a part of. The VPN configurations along with the other device template policies and configurations enable easy deployment of the branch device to your network. Device templates support the configuration of a device only as a spoke. A device can be part of multiple hub and spoke site-to-site VPN topologies.

After the configured device template is applied to a device, the variables are resolved, the protected network overrides are configured, and the device is added as a spoke in the specified VPN topology.

Variables and network object overrides

You can parameterize template configurations using variables and network object overrides.

Variables

A variable is an object type that is supported for template configurations. A variable in a template defines specific configuration values for a device. You can define values for these variables during application of the template on the device. You can see the variable icon (x) for the fields that use a variable. The variables are displayed with a $ prefix to distinguish these values from the other values.

For information on supported variable types and creating variables, see Supported Variables and Add a Variable.

Network object overrides

Network object overrides are similar to variables. But, these are used to provide override values for a network object. You can declare a list of network objects in the template and create network object overrides for these objects. You can then provide values for these network object overrides during the application of the template on the device. For example, if you define a host network object in a template, you can add a network object override before the application of the template on the device and then provide a relevant value during the application of the template on the device.

For more information on supported network objects and adding a network object override, see Supported Network Object Overrides and Add a Network Object Override.

Model mapping

Model mapping lets you map template-defined interfaces to the interfaces of the Firewall Threat Defense model you require.

Model mapping requirements

Create model mappings in the template before applying the template to the device. During template application, the variables in the interface configurations are replaced with the values that you provide, then the resulting configurations are copied to the mapped interfaces on the device.

Interface configurations differ for each device model, so you must copy interface configurations from the template to the device's target interfaces.

During template application, the variables in the interface configurations are replaced with the values that you provide, then the resulting configurations are copied to the mapped interfaces on the device.

For more information on setting up model mapping, see Add Model Mapping.

Templates and high availability

You can apply device templates on Firewall Threat Defense high availability devices after device registration. High availability-specific configurations are not supported in device templates. Any high availability configurations and monitored interfaces that are already part of the target high availability device pair configurations are not modified. You cannot map any template interfaces to failover interfaces.

You can generate a device template from a high-availability device pair. Template operations such as application of the template on a device, template generation, importing, and exporting of a template, can be performed only on the active unit. You cannot perform these operations on the standby unit.

Validating template configurations before and after application on devices

Summary

The key components involved in the template configuration validation process are:

- System validator: Performs automated checks to ensure device compatibility and configuration integrity
- Target device: The network device receiving the template configuration
- Template configuration: The configuration settings being applied to the device
- Rollback mechanism: System capability to restore original device state if validation fails

Workflow

These stages describe how template configuration validation works:

1. The system performs pre-application validation checks before deploying the template on the target device.
   - Ensure that the target device model and version are supported.
   - Cluster and container checks -The device must not be part of a cluster or multi-instance.
   - Model mapping validation - Model mapping for the target device model exists and is valid.
   - Sanity check of template parameter values. For example, two variables used as IP addresses of interfaces must not have the same value.
2. The system applies the template configuration to the target device after successful pre-application validation.
3. The system performs post-application validation checks to ensure the applied configurations are valid.
   - Interface configuration validation. For example, variables used for the IP address fields of two or more interfaces must not have the same IP address values.
   - Routing policy validation. For example, the IPv4 address in BGP neighbor configurations must not overlap with the IP address of any interface.

   If the validation checks that are done at the end of the task to apply the template on the device fail, any applied configurations are rolled back and the device is restored to it's original state.

Prerequisites for device registration using device templates

This reference provides the prerequisites for device registration using device templates, including supported models and VPN configuration requirements.

Model support

Device templates are supported on On-Prem Cloud-Delivered Firewall Management Center, Cloud-Delivered Firewall Management Center(cdFMC), with these models running Secure Firewall version 7.4.1 and later versions:

- Secure Firewall 200

- Firepower 1010
- Firepower 1100
- Secure Firewall 1200
- Firepower 2100 (7.4.x only)
- Secure Firewall 3100

Prerequisites for VPN connections in device templates

- Configure site-to-site VPN topologies that must be used in the device template.
- Ensure that you have configured all hub and VPN topology-related configurations such as authentication methods, IKE and IPsec policies.
- Supported types of VPN hub and spoke topologies are:
  - Policy-based
  - Route-based
  - SD-WAN
- Assign appropriate logical names and IP addresses to the interfaces of the Firewall Threat Defense devices. For example, use *inside* for the interface connected to the LAN, and *outside* for the interface connected to the internet or WAN.
- Spoke devices must be version 7.4.1 and later.

Licenses for device registration using device templates

This reference provides license requirements and conditions for device registration using device templates, including Smart Licensing account requirements and Secure Client license behavior during template application.

- Device templates do not have any specific license requirements.
- License entitlements for the target device must be present in the Smart Licensing account.
- To configure VPN connections in the template, the Essentials license must allow export-controlled functionality. Choose **System** (⚙) > **Licenses** > **Smart Licenses** to verify this functionality in the Cloud-Delivered Firewall Management Center.
- When you apply a template on a device, note these conditions for Secure Client licensing:

Device with Secure Client License	Template with Secure Client License	Secure Client License after Device Template Application
Yes	Yes	Template License
Yes	No	Device License
No	Yes	Template License

Guidelines for device registration using device templates

General guidelines

Follow these general guidelines when using device templates for device registration and configuration management.

- All device configurations other than VNI and VTEP are supported.
- You can attach shared policies and S2S VPN policies to a template. These policies are assigned during template application.
- Templates can be applied on HA devices. However, application of device templates during HA device pair registration is not supported. You also cannot manage HA-related configurations such as failover links, standby IP addresses, and so on. For more information, see Device Template Operations on Threat Defense HA Devices.
- Device template operations are supported only on the active Cloud-Delivered Firewall Management Center. The standby peer does not support device template operations.
- Ensure that the template names and device display names are not the same.
- Ensure that you do not create or delete a template during device backup or restore operations.
- After application of the template on the device, if the manager access is changed from management to data interface or vice-versa, you must re-establish the management connection with the device. Note that you cannot change the manager access interface during template application.
- You can add a maximum of 250 device templates to the Cloud-Delivered Firewall Management Center.
- Templates that are created and configured for devices that are managed through the data interface cannot be used to register and apply to devices that are managed through the management interface.
- Device registration and application of template does not come under the Change Management workflow. Only approved data, such as access policies, templates, template variables, network overrides declared in template, and template configurations used in the template application operation are used.
- Device registration with serial number and access control policy is supported for only one device at a time.
- When you apply a template to a device that is already registered, the configuration in the template is only copied to the target device. You can then choose to manually deploy the configuration on the device or let the copied configuration stay on the device and deploy later. However, if you apply the template during device onboarding, the configuration in the template is copied to the target device and, as per existing behavior, automatically deployed on the device after device registration.
- Any change in model mapping causes the device to be marked as 'Out of Sync'. Consider reapplying the template to the devices if you have made changes to the interface mapping in the corresponding model mapping or if you have made any configuration changes after the previous application of the template.
- Device templates are only supported on merged management and diagnostic interfaces. For more information, see Merge the Management and Diagnostic Interfaces.
- Template configuration updates are supported by change management. Creation and application of device templates is not supported by change management.

- You cannot sync configuration changes from a device to the template. A few sample scenarios during which you may want to make changes to the device configuration using a template along with recommended solutions are given below:
  - If you want to test the new configuration changes on one device before propagating the changes to multiple devices, we recommend that you make the changes in the template and apply the template to one device. Validate the changes on that device and then apply the template to the other devices.
  - If you want to make a large number of changes resulting in a signification deviation from the current configuration on a device, and then propagate those changes to other devices, you may choose one of these options:
    - Export the current template to get a copy of the template. You can then make the required changes in the template and apply to a single device. Validate the changes on that device and then apply the template to the other devices.
    - You can also make the required changes on the device and create a template from that device. You can then apply and validate the changes on the other devices. However, we do not recommend this as template parameters such as variables and network object overrides will not be present in the created template.
  - If the configuration on a particular device starts to differ significantly from the configuration in the template, you may also choose to not use the template for this device and delete the device-template association from the **Associated Devices** window.

Guidelines for VPN connections

Follow these guidelines when configuring VPN connections in device templates.

- Supported interfaces for VPN topologies are:

Topology Type	Interface Type

Policy-Based and SD-WAN	• Physical interfaces
	  • Non-management
	  • Interface Mode must be either Routed or None
	• Subinterfaces
	• Redundant interfaces
	• Etherchannel interfaces
	• VLAN interfaces

Route-Based	Static Virtual Tunnel Interfaces

- When you apply a template on a device that is part of a VPN topology, you must ensure that the template includes interface configurations for all interfaces used in the topology.

- When you apply a template with VPN connections to multiple devices, note that a template is applied to multiple devices in the order in which you have selected the devices. If the template has VPN connections, the corresponding VPN topology is locked.
- For SD-WAN topology VPN connections: Ensure that IP address subnet of the interface does not conflict with the subnet of the IP address pool of the SD-WAN hub.
- Change management: Before you apply a device template to a device, ensure that the VPN topology is not locked by a Change management ticket.

Unsupported configurations

The following features and configurations are not supported using device templates.

- Multi-instance mode
- Clustering
- Non-converged management interface
- Transparent mode
- HA failover configurations
- Chassis configurations
- Logical devices
- Variables for nested objects
- Override support for network groups and other object types

Limitations for VPN connections

Be aware of these limitations when working with VPN connections in device templates.

- When you create a template from a device that is part of a VPN topology (**Devices** > **Device Management** **More** (⋮) **> Generate Template from Device**), VPN configurations are not part of the template. You must reconfigure the VPN configurations on the template.
- When you export a device template with one or more VPN connections (**Template Settings > General > General** pane > **Export**), the VPN connections are not exported. You must reconfigure the VPN connections on the imported template.
- Certificate-based authentication:
  - Device templates do not support automatic certificate enrolment of a device.
  - When you onboard a device using a template with VPN configurations, if the VPN topology uses certificate-based authentication, the first deployment to the device will fail. Ensure that you manually enroll the device certificate after the device registration and deploy the configurations on the device again.

Configure device templates

Add a template and then configure settings in the template.

Add a device template

You can add a new device template with the required configuration or generate a device template from an existing device.

Create a new device template

Create a device template to establish standardized configurations that can be applied to multiple devices, ensuring consistent security policies and routing modes across your network.

You can specify the device template name, description, access control policy, and routing mode. You can add more configurations after the template has been created.

Follow these steps to create a device template:

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click **Add Device Template**.

Step 3 In the **Add Device Template** window, enter a **Name** for the template.

Step 4 (Optional) Enter a **Description** for the template.

Step 5 Choose an **Access Control Policy** from the drop-down list.

Step 6 Choose a **Mode** from the drop-down list.

Step 7 Click **OK**.

Generate a new device template from an existing device

The new template has the same configuration as the device from which it is generated. You can generate a new device template from standalone and HA devices. However, if you generate a template from HA devices, the new template will not contain the failover configurations.

You can generate a new device template from a device that is registered with the Cloud-Delivered Firewall Management Center.

Follow these steps to generate a new device template from an existing device:

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the **More** (⋮) icon, and click **Generate Template from Device**.

In the new device management view, select the device, and then select **Generate template from device** in the **Device actions** pane.

Step 3 In the **Generate template from device** window, enter a **Name** for the template.

Step 4 (Optional) Enter a **Description** for the template.

Step 5 Choose an **Access Control Policy** from the drop-down list.

Step 6 Click **OK**. To view the status of template creation, go to **Notifications**, and then click the **Tasks**tab.

Step 7 Choose **Devices** > **Template Management** to view the newly created template.

Import a device template

The device template import feature allows you to transfer templates between different Cloud-Delivered Firewall Management Center systems or domains for replication, modification, and deployment purposes.

You can import a template into the Cloud-Delivered Firewall Management Center or export a template to your local system. This feature is useful in the following scenarios:

- Generate a copy of the template from a device, export it, and import that template into another Cloud-Delivered Firewall Management Center or Cloud-Delivered Firewall Management Center.
- Generate a copy of the template, export it, modify as required to create a variation of the existing template, and import the template into the Cloud-Delivered Firewall Management Center.
- Generate a copy of the template, export it, and import that template into another domain in which the source template is not visible.

Follow these steps to import a device template into the Cloud-Delivered Firewall Management Center from your local system:

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **More** (⋮) icon for the template that you want to replace with an imported template.

Step 3 Click **Import** and click **Import** again.

If you want to export a template, click **Export** and click **OK**.

Step 4 To view the status of the import task, click **Notifications**, and then click the **Tasks** tab.

Step 5 Choose the template SFO file on your local system and click **Open**. This template SFO file that you import can be newly created, generated from a device, or cloned from an existing template.

Step 6 To view the status of the import task, click **Notifications**, and then click the **Tasks** tab. You will see a notification informing you that the import or export task is successfully completed. If you are exporting a template, click **Notifications** and click the **Tasks**tab. In the tab page, click **Download Export Package** to download the template configuration as an SFO file.

Note

Alternatively, you can also go to the **Template Management** window and click the **Edit** (✎) icon of the template. Then, go to **Template Settings**, and click **General**. Click **Import** or **Export** in the **General** tile to import or export the template.

Configure device settings in the template

After creating the template, you can set up device configurations and adjust the settings by editing the template.

Add a physical interface

A physical interface allows you to expand network connectivity in a device template beyond the default interfaces provided.

By default, a device template will enable the device to come up with the following physical interfaces:

- Management interface
- Inside interface
- Outside interface

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template in which you want to add the physical interface.

Step 3 In the **Interfaces** tab, click **Add Physical Interface**.

Step 4 Choose a **Slot** and **Port Index** number from the drop-down list.

Step 5 Click **Create Interface**.

Add a logical interface

Create a logical interface in a template to define network connectivity and traffic handling for your firewall deployment.

You can create a logical interface in the same way as you do on the Cloud-Delivered Firewall Management Center without using the template. Perform the procedure given below to create a logical interface.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template in which you want to add the logical interface.

Step 3 In the **Interfaces** tab, click **Add Interface**, and choose the type of interface that you want to create from the drop-down list. You can create the following types of interfaces:

- Sub-interface
- Ether channel interface
- Bridge group interface
- VLAN interface
- Virtual tunnel interface
- Loopback interface

For more information, see Interface overview and Regular Firewall Interfaces.

Edit an interface

Edit interface configurations on Firepower 1000,Secure Firewall 1200, Firepower 2100, and Secure Firewall 3100 Firewall Threat Defense devices using device templates.

You can edit an interface in the same way as you do on the Cloud-Delivered Firewall Management Center without using the template. Use template variables to set up the IPv4 and IPv6 addresses. The device template supports the configurations that are supported on Firepower 1000,Secure Firewall 1200, Firepower 2100, and Secure Firewall 3100 Firewall Threat Defense devices.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template in which you want to edit the physical interface.

Step 3 In the **Interfaces** tab, click the **Edit** icon for the interface that you want to edit.

Step 4 In the **Edit Physical Interface** window, you can edit any of the following settings:

- General
- PoE
- IPv4
- IPv6
- Path Monitoring
- Hardware Configuration
- Manager Access
- Advanced

Note
Use variables to configure IPv4 and IPv6 addresses. For more information on templatizing variables, see Configure Template Parameters.

For more information on editing the settings mentioned above, see Interface overview and Regular Firewall Interfaces.

Configure other device settings

Configure additional device settings within templates to maintain consistent configuration across multiple devices.

Configure the other device settings in the same way as you do on the Cloud-Delivered Firewall Management Center without using the template.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon for the template in which you want to configure the settings.

Step 3 Click the tabs at the top of the window to configure any of these settings:

- Inline Sets
- Routing
- DHCP
- VPN
- Template Settings

Template settings

When you apply a template to a device, the template-specific settings are copied to the device.

Available template settings

you can configure these template settings in the **Template Settings** window.

- General
  - General
  - License
  - Applied Policies
  - Advanced Settings
  - Deployment Settings
- Template Parameters

- Variables
- Network object overrides

- Model Mapping

Edit general settings

Edit the name of the device template and enable or disable packet transfer to customize the device configuration according to your requirements.

In the **General** tile, you see the following fields:

- **Template Name** – The name of the template.
- **Transfer Packets** – Displays whether or not the managed device sends packet data with the events to the management center.
- **Mode** – Displays the mode of the management interface for the device: routed.
- **Configuration** – Click **Export** to export the template configurations as an SFO file. Click **Import** to import an SFO file that has the template configurations that you require.
- **Manage device by Data Interface** – Toggle the button to enable or disable management of the device using the data interface.

Follow these steps to edit the name of the device and to enable or disable packet transfer:

Procedure

Step 1 Click the **Edit** (✎) icon in the **General** tile.

Step 2 Change the **Template Name** as per your requirement.

Step 3 Check the **Transfer Packets** checkbox to allow packet data to be stored with events on the Cloud-Delivered Firewall Management Center.

Step 4 Click **Save**.

Edit licenses

Edit license types to match your requirements for template configurations.

In the **License** tile, you can see the **License types** that are required based on the configurations used in the template. Choosing a license here does not consume that license on the device. The license is consumed only when you apply the template to a device.

Procedure

Step 1 Click the **Edit** (✎) icon in the **License** tile.

Step 2 Check or clear the check box next to the license you want to enable or disable for the managed device.

Step 3 Click **Save**.

Edit applied policies

Edit applied policies to modify the access control policy assignments for a template according to your requirements.

In the **Applied Policies** tile, you can see the access control policies that are associated with the template.

For policies with links, you can click the link to view the policy.

Procedure

Step 1 Click the **Edit** (✎) icon in the **Applied Policies** tile.

Step 2 For each policy type, choose a policy from the drop-down list. Only existing policies are listed.

Step 3 Click **Save**.

Edit advanced settings

Edit advanced configuration settings to optimize device performance and memory usage according to your network requirements.

The **Advanced Settings** tile displays the advanced configuration settings, as described below. You can edit any of these settings.

Table 2: Advanced section table fields

Field	Description
Application Bypass	The state of Automatic Application Bypass on the device.
Bypass Threshold	The Automatic Application Bypass threshold, in milliseconds.
Object Group Search	The state of object group search on the device. While operating, the Firewall Threat Defense device expands access control rules into multiple access control list entries based on the contents of any network or interface objects used in the access rule. You can reduce the memory required to search access control rules by enabling object group search. With object group search enabled, the system does not expand network or interface objects, but instead searches access rules for matches based on those group definitions. Object group search does not impact how your access rules are defined or how they appear in the Cloud-Delivered Firewall Management Center. It impacts only how the device interprets and processes them while matching connections to access control rules. **Note** By default, the **Object Group Search** is enabled when you add Firewall Threat Defense for the first time in the Cloud-Delivered Firewall Management Center.

Field	Description
Interface Object Optimization	The state of interface object optimization on the device. During deployment, interface groups and security zones used in the access control and prefilter policies generate separate rules for each source/destination interface pair. If you enable interface object optimization, the system will instead deploy a single rule per access control/prefilter rule, which can simplify the device configuration and improve deployment performance. If you select this option, also select the **Object Group Search** option to reduce memory usage on the device.

Procedure

Step 1 Click the **Edit** (✎) icon in the **Advanced Settings** tile.

Step 2 Change the settings as per your requirement.

- Configure Automatic Application Bypass
- Configure Object Group Search
- Configure Interface Object Optimization

Step 3 Click **Save**.

Edit deployment settings

Edit deployment settings to configure auto rollback functionality when management connection fails due to misconfiguration and to set the time interval for connectivity monitoring.

The **Deployment Settings** tile displays the information described in the table below.

Table 3: Deployment settings

Field	Description
Auto Rollback Deployment if Connectivity Fails	Enabled or Disabled. You can enable auto rollback if the management connection fails as a result of the deployment; specifically if you use data for Cloud-Delivered Firewall Management Center access, and then you misconfigure the data interface.
Connectivity Monitor Interval (in Minutes)	Shows the amount of time to wait before rolling back the configuration.

Deployment settings include enabling auto rollback of the deployment if the management connection fails as a result of the deployment; specifically if you use data for Cloud-Delivered Firewall Management Center access, and then you misconfigure the data interface. Alternatively, manually roll back the configuration using the **configure policy rollback** command.

Procedure

Step 1 Click the **Edit** icon in the **Deployment Settings** tile.

Step 2 Set the **Connectivity Monitor Interval (in Minutes)** to set the amount of time to wait before rolling back the configuration. The default is 20 minutes.

Step 3 If a rollback occurs, see the following for next steps.

- If the auto rollback was successful, you see a success message instructing you to do a full deployment.
- You can also click **Deploy** and then go to the **Advanced Deploy** screen and click the **Preview** icon to view the parts of the configuration that were rolled back (see Deploy Configuration Changes). Click **Show Rollback Changes** to view the changes, and **Hide Rollback Changes** to hide the changes.
- In the Deployment History Preview, you can view the rollback changes.

Step 4 Check that the management connection was reestablished.

In Cloud-Delivered Firewall Management Center, check the management connection status on the Connection Status page. Navigate to **Devices** > **Device Management** page and then navigate to **Management** area under the **Device** tab. Then in the **FMC Access Details** screen, click **Connection Status**.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface.

Configure template parameters

You can templatize configurations using template parameters such as variables and network object overrides.

Supported variables

The variable types are supported in device templates.

Variable Name	Description	Type
AS Number	Defines the unique Autonomous System (AS) number.	Integer Example: 2
FQDN	Defines a single Fully Qualified Domain Name (FQDN).	String Example: abc.example.com
IPv4 Host	Defines the IPv4 address of the host.	String Example: 209.165.201.8
IPv4 Network	Defines the IPv4 network address block.	String Example: 209.165.200.224/27

Variable Name	Description	Type
IPv4 Range	Defines the range of IPv4 addresses.	String Example: 209.165.200.225-209.165.200.250
IPv6 Host	Defines the IPv6 address of the host.	String Example: 2001:DB8::1
IPv6 Network	Defines the IPv6 network address block.	String Example: 2001:DB8:0:CD30::/60
Password	Defines a password string.	String Example: E28@2OiUrhx!
Router ID	Defines an identifier for the router.	Integer Example: 21
String	Defines a custom string.	String Example: testvalue2

Add a variable

Add variables to create reusable configuration elements that can be applied across multiple device templates, enabling consistent and efficient policy management.

Variables provide a way to standardize configuration elements across your deployment. You can create variables for network objects, ports, or other configurable elements that need to be consistent across multiple devices or policies.

Follow these steps to add a variable:

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Variable** from the list of object types.

Step 3 Click **Add Variable**.

Step 4 Enter a **Name**.

Step 5 Choose a **Variable Type** from the drop-down list.

Step 6 (Optional) Enter a **Description**.

Step 7 Click **Save**.

Supported network object overrides

This reference provides information about the network objects that are supported for override configurations, including their descriptions and data types with examples.

Network Object Name	Description	Type
Network	An address block, also known as a subnet.	String Example: IPv4 - 209.165.200.224/27 IPv6 - 2001:DB8::/48
Host	The IP address of the host.	String Example: IPv4 - 209.165.200.225 IPv6 - 2001:DB8:1::1
Range	A range of IP addresses.	String Example: IPv4 - 209.165.200.225-209.165.200.250 IPv6 - 2001:DB8::1 - 2001:DB8:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF
FQDN	A single fully-qualified domain name (FQDN).	String Example: abc.example.com

Add a network object override

Network object overrides allow you to customize network objects within templates for specific deployments without modifying the base template.

Use this procedure to add network object overrides to a template when you need to customize network objects for specific deployment scenarios.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template in which you want to add the network object override.

Step 3 Choose **Template Settings** > **Template Parameters**.

Step 4 In the **Network Object Overrides** section, click **Add or Remove Network Object Overrides**.

Step 5 In the **Add or Remove Network Object Overrides** window, choose the network objects for which you want to create network object overrides from the **Available Networks** window and click the > button.

Step 6 Click **Save**.

Add model mapping

Model mapping allows you to map a template to one or more models as long as the interface configurations are valid for all mapped models.

For each model, you can specify which template interface corresponds with which model interface. For example, if the template includes switch ports and VLAN interfaces, then that template can only be applied to a Firepower 1010 or Secure Firewall 1210/1220.

Follow these steps to add model mapping:

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click **Add Model Mapping** for the template in which you want to create the model mapping. Alternatively, you can click the **Edit** (✎) icon of the template and choose **Template Settings** > **Model Mapping**.

Step 3 Click **Add Model Mapping**.

Step 4 Choose the **Device Model** from the drop-down list.

Step 5 Map the template interfaces to the device model interfaces by choosing the interface from the **Model Interface** drop-down list.

Note
You can click **Clear Mapping** to remove the defined model mapping. Click **Reset Mappings** for default interface mapping in which the mapping is done based on the slot and port index order of the interface names.

Step 6 Click **Save**. The interface mappings are listed along with the device model and mapping status on the **Model Mapping** window.

Note
Some configurations in the template may not be supported on all device models. Unsupported configurations, if any, are not applied to the device. The device Template Apply Report provides details about such configurations.

Invalid model mappings

Some configurations in the template may not be supported on all device models. Unsupported configurations, if any, are not applied to the device. Valid model mappings can become invalid when you modify template configurations. For example, if you add a new interface on the template and assign a name to it, you must map the new interface to the appropriate device model interface.

Causes of invalid model mappings

Model mapping can be invalidated due to any of these reasons:

- Number of configured VRF instances exceeds the limit for a specific model.
- Interfaces mapped to incompatible models, versions, or interfaces. See Requirements and Prerequisites for more information.
- Number of interfaces exceeds the model limit.

- Deleted an interface that was mapped.
- Newly added physical interfaces are not mapped to a compatible model interface.
- Model mapping is not done for a named interface.
- Model mapping is not done for an interface related to other logical interfaces, such as sub-interfaces, PC interfaces, and so on.
- Policy or configuration changes that are unsupported on some device models, such as enabling switch port configuration on interfaces.

You can also save a template with invalid model mappings. However, you must review and fix the model mapping before initiating application of the template on the device.

You can hover over **Invalid** under **Mapping Status** to view the errors that caused the invalid mapping status. Fix the errors before initiating application of the template on the device.

Configure Site-to-Site VPN connections in a device template

This reference provides the configuration parameters and settings required to establish Site-to-Site VPN connections using device templates.

Configure an SD-WAN VPN connection

This task allows you to add spokes to SD-WAN topologies by configuring an SD-WAN VPN connection through the device template.

You can configure an SD-WAN VPN connection to add spokes to SD-WAN topologies using the device template.

Before you begin

- Configure a minimum of one SD-WAN topology (**Devices** > **VPN** > **Site to Site**).
- Review Prerequisites for Configuring Device Templates and Guidelines and Limitations for Device Templates.

Follow these steps to configure an SD-WAN VPN connection:

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the edit icon adjacent to the device template that you want to edit.

Step 3 Click the **VPN** tab.

Step 4 Click **Add VPN Connection**.

Step 5 Choose an SD-WAN topology from the **VPN Topology** drop-down list.

The **Add VPN Connection** dialog box expands and you can configure the following parameters:

a) From the **VPN Interface** drop-down list, choose a WAN-facing or internet-facing physical interface to establish a VPN connection with the hub.

This list contains all the interfaces configured on the device template.

b) **Use IP Address from the VPN Interface**—This drop-down list is auto populated with the IP address variable. For IPv6 addresses, choose an IPv6 address from the drop-down list.
c) Check the **Local Tunnel (IKE) Identity** check box to enable a unique and configurable identity for the VPN tunnel from the spoke to a remote peer.
d) **Identity Type**—Key ID is the only supported identity type. Choose a key ID variable from the drop-down list or click (+) to create a new key ID variable.
e) Click **OK**.

You can view the VPN connection in the **Site-to-Site VPN Connections** table.

Step 6 Click **Save**.

What to do next

1. Configure the routing policy for the spoke in the device template.
2. Map the device interfaces to the template interfaces (Model Mapping).
3. Apply the template to a device.

Configure a route-based Site-to-Site VPN connection

This task allows you to add spokes to existing route-based site-to-site VPN topologies by configuring VPN connections through the device template interface.

You can configure a route-based site-to-site VPN connection to add spokes to route-based site-to-site VPN topologies using the device template.

Before you begin

- Configure a minimum of one route-based site-to-site VPN topology (**Devices** > **VPN** > **Site to Site**).
- Review Prerequisites for Configuring Device Templates and Guidelines and Limitations for Device Templates.

Follow these steps to configure a route-based site-to-site VPN connection:

Procedure

Step 1 Choose **Devices** > **Template Management**.
Step 2 Click the edit icon adjacent to the device template that you want to edit.
Step 3 Click the **VPN** tab.
Step 4 Click **Add VPN Connection**.
Step 5 Choose a route-based site-to-site VPN topology from the **VPN Topology** drop-down list.

The **Add VPN Connection** dialog box expands and you can configure the following parameters:

a) From the **Virtual Tunnel Interface (VTI)** drop-down list, choose a VTI interface or click (+) to create a new VTI.

VTI is a virtual interface used to establish a route-based VPN tunnel. You must configure routing policies for a VTI to set up a VPN tunnel. This list contains all the VTIs configured on the device template. For more information on creating a VTI, see Add a VTI interface, on page 1752.

b) Check the **Use Public IP Address** check box to override the tunnel source IP address and configure a public IP address variable for the VTI. Click (+) to create a new public IP address variable.

This IP address is the source IP address for the VPN tunnel. By default, this is the IP address of the VPN interface. However, if the device is behind NAT, the VPN interface has a private address, but the post-NAT public IP address should be configured.

c) Check the **Local Tunnel (IKE) Identity** check box to enable a unique and configurable identity for the VPN tunnel from the spoke to a remote peer.

d) **Identity Type**: Key ID is the only supported identity type. Choose a key ID variable from the drop-down list or click (+) to create a new key ID variable.

e) (Optional) Check the **Enable Secondary VPN Tunnel** check box to configure the parameters for the secondary VPN tunnel.

f) Click **OK**.

You can view the VPN connection in the **Site-to-Site VPN Connections** table.

Step 6 Click **Save**.

What to do next

1. Configure the routing policy for the spoke in the device template.
2. Map the device interfaces to the template interfaces (Model Mapping).
3. Apply the template to a device.

Configure a policy-based Site-to-Site VPN connection

This task enables you to add spokes to policy-based site-to-site VPN topologies using the device template.

You can configure a policy-based site-to-site VPN connection to add spokes to policy-based site-to-site VPN topologies using the device template.

Before you begin

- Configure a minimum of one policy-based site-to-site VPN (**Devices** > **VPN** > **Site to Site**).
- Review Prerequisites for Configuring Device Templates and Guidelines and Limitations for Device Templates.

Follow these steps to configure a policy-based site-to-site VPN connection:

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the edit icon adjacent to the device template that you want to edit.

Step 3 Click the **VPN** tab.

Step 4 Click **Add VPN Connection**.

Step 5 Choose a policy-based site-to-site VPN topology from the **VPN Topology** drop-down list.

The **Add VPN Connection** dialog box expands and you can configure these parameters:

a) From the **VPN Interface** drop-down list, choose a WAN-facing or internet-facing physical interface to establish a VPN connection with the hub.

This list contains all the interfaces configured on the device template.

Do one of these to configure the IP address of the VPN interface:

- Click the **Use IP Address from the VPN Interface** radio button to use the IP address of the VPN interface.

  This IP address is auto populated. For IPv6 addresses, choose an IPv6 address from the drop-down list.

- Click the **Use Public IP Address** radio button to configure a public IP address for the VPN interface.

  Choose an IP address variable from the drop-down list or click (+) to add an IP address variable.

b) Check the **Local Tunnel (IKE) Identity** check box to enable a unique and configurable identity for the VPN tunnel from the spoke to a remote peer.

c) **Identity Type**: Key ID is the only supported identity type. Choose a key ID variable from the drop-down list or click (+) to add a new key ID variable.

d) **Protected Networks**: Click (+) to configure a protected network for the VPN connection.

Do one of these:

- Choose a protected network and click **OK**.
- Click **Add** to configure a network object and click **Save**.

  When you create a protected network object, note this:

  - Click either the **Host** or the **Network** radio button.
  - Check the **Allow Overrides** check box.

e) Click **OK**.

You can view the VPN connection in the **Site-to-Site VPN Connections** table.

Step 6 Click **Save**.

What to do next

1. Before you apply a template to a device, to configure device-specific values for the protected networks, add these objects in **Template Settings** > **Template Parameters** > **Add Network Objects Overrides**.
2. Map the device interfaces to the template interfaces (model mapping).
3. Apply the template to a device.

Add a device to an SD-WAN topology in a dual ISP deployment

This task enables you to add a device to an SD-WAN topology in a dual ISP deployment using a device template.

Use this procedure to integrate a device into an existing SD-WAN topology that utilizes dual ISP connections for redundancy and load balancing.

Before you begin

Ensure that you have two SD-WAN VPN topologies with the same hub. For more information about configuring SD-WAN topologies, see Configure an SD-WAN topology using the SD-WAN wizard, on page 1792.

Follow these steps to add a device to an SD-WAN topology in a dual ISP deployment:

Procedure

Step 1 Create a device template.

Navigate to **Devices** > **Template Management**.

For detailed instructions, see Create a new device template, on page 126.

Step 2 Add a physical interface in the template.

By default, a template has only one outside interface. Rename the outside interfaces, for example, ISP1, ISP2.

Navigate to **Devices** > **Template Management**, and click **Interfaces > Add Physical Interface** .

For detailed instructions, see Add a physical interface, on page 128.

Step 3 Configure an SD-WAN VPN connection using ISP1 interface.

Navigate to **Devices** > **Template Management**, and click **VPN > Add VPN Connection**.

For detailed instructions, see Configure an SD-WAN VPN connection, on page 138.

Step 4 Configure an SD-WAN VPN connection using ISP2 interface.

Navigate to **Devices** > **Template Management**, and click **VPN > Add VPN Connection**.

For detailed instructions, see Configure an SD-WAN VPN connection, on page 138.

Step 5 Add static routes from ISP1 and ISP2 interfaces to the SD-WAN hub network.

Navigate to **Devices** > **Template Management**, and click **Routing > Static Route**.

Step 6 Add the ISP1 and ISP2 interfaces to an ECMP zone.

Navigate to **Devices** > **Template Management**, and click **Routing > ECMP**.

Step 7 Configure the network object overrides.

Navigate to **Devices** > **Template Management**, and click **Template Settings > Template Parameters > Add Network Objects Overrides**.

For detailed instructions, see Add a network object override, on page 136.

Step 8 Map the template interfaces to the device model interfaces (Model Mapping).

Navigate to **Devices** > **Template Management**, and click **Template Settings > Model Mapping**.

For detailed instructions, see Add model mapping, on page 137.

Step 9 Apply template to the device.

Navigate to **Devices** > **Template Management**.

For detailed instructions, see Apply a template, on page 145.

The device is successfully added to the SD-WAN topology in a dual ISP deployment. For more information about dual ISP deployment using the SD-WAN wizard, see Sample configurations for dual-ISP deployment using SD-WAN wizard, on page 1797.

What to do next

Click **Deploy** deploy the configurations on the device and the hubs of the SD-WAN topologies.

Configure a Template for Firewall Threat Defense Devices Managed Using the Data Interface

This task configures a template for Firewall Threat Defense devices that are managed using a data interface, ensuring proper connectivity parameters match between the device and template to maintain Cloud-Delivered Firewall Management Center connectivity after template deployment.

To configure a template for a Firewall Threat Defense device that is managed using a data interface for Cloud-Delivered Firewall Management Center connectivity, ensure that the connectivity parameters of the device match the template. This ensures that the Threat Defense device does not lose connectivity with the Cloud-Delivered Firewall Management Center after application of the template. A template that you configure for Firewall Threat Defense devices managed using the data interface cannot be applied on devices that are not managed by the data interface.

These are the connectivity parameters:

- Data interface used to the manage the Firewall Threat Defense device. For example, **Ethernet1/1**.
- Name of the interface. For example, **outside**.
- IP address configured on the data interface. For example, DHCP or static IP.
- Route configured for the data interface. This can be a default or specific route defined on the data interface used for connectivity between the Firewall Threat Defense device and the Cloud-Delivered Firewall Management Center.
- DDNS hostname configuration on the data interface.

If the template connectivity parameters do not match those on the device, validation fails and the template is not applied on the device. Validation does not require an exact match for some parameters, such as IP address or DDNS hostname. However, configure these parameters appropriately to maintain connectivity between the Firewall Threat Defense device and the Cloud-Delivered Firewall Management Center after deployment.

These template validation checks are done to ensure sanity of configurations that are required to manage the Firewall Threat Defense device using the data interface:

- You cannot apply a template in which manager access to the device is configured with the management interface to a device in which manager access to the device is configured with the data interface.
- You cannot apply a template in which manager access to the device is configured with the data interface to a device in which manager access to the device is configured with the management interface.
- You cannot apply a template in which manager access to the device is configured with the single WAN data interface to a device in which manager access to the device is configured with the dual WAN data interface.
- If any of the connectivity parameters do not match, you cannot apply a template in which manager access to the device is configured with the data interface to a device in which manager access to the device is configured with the data interface.

Before you begin

Follow these steps to configure a template for Firewall Threat Defense devices managed using the data interface:

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template that you want to configure to manage Firewall Threat Defense devices using the data interface.

Step 3 Click the **Template Settings** tab.

a) In the **General** tile, toggle the **Manage device by Data Interface** button.
b) You will see a popup asking you to pick a data interface for manager access. Click **OK**.

Step 4 Click the **Interfaces** tab.

a) Click the **Edit** icon of the data interface that you want to use for manager access. The first data interface – Ethernet1/1 (outside interface), is the data interface that is most commonly used for manager access.
b) In the **Edit Physical Interface** window, click the **Manager Access** tab.
c) Check the **Enable management access** checkbox.
d) Click **OK**. You will see that the interface that you selected for manager access has been marked with **Manager Access**.

Step 5 Click the **DHCP** tab.

Step 6 Click the **DDNS Update Methods** tab.

a) Click **+Add** to add a DDNS update method.
b) In the **Add DDNS Update Method** window, enter a **Method Name** and choose **FMC only**.
c) Set the **Update Interval** as per your requirement.
d) Click **OK**. You will see the method that you created in the **DDNS Update Methods** table.

Step 7 Click the **DDNS Interface Settings** tab.

a) Click **+Add** to add dynamic DNS configuration.

b) In the **Add Dynamic DNS** configuration window, choose values for these fields:

- **Interface** – Choose the interface enabled for manager access
- **Method Name** – Choose the method that you created.
- **Host Name** – Choose a variable for the hostname.

Do not edit the rest of the fields in this window.

c) Click **OK**. The **DDNS Interface Settings** table is populated with the entry that you created.

Step 8 To configure the model mapping to ensure that the data interface set for manager access in the template matches the data interface selected for manager access on the device, click the **Template Settings** tab and click **Model Mapping**.

a) Click **Add Model Mapping**.

b) Choose the **Device Model** from the drop-down list.

c) Map the date interface that is set for manager access in the template to the appropriate data interface on the device by choosing the interface from the **Model Interface** drop-down list.

d) Click **Save**. The interface mappings are listed along with the device model and mapping status on the **Model Mapping** window. You can now apply the template on a device that is managed using the data interface.

Use templates with devices

You can apply templates in these scenarios:

- During device registration: Choose a compatible template for the initial configuration when registering a device with the Cloud-Delivered Firewall Management Center.
- After device registration: Apply templates to devices that are already registered with the Cloud-Delivered Firewall Management Center.

Apply templates to existing devices

You can apply or reapply a template to existing devices.

Apply a template

Template application clears existing configurations and applies configurations from the template to devices already registered with the Cloud-Delivered Firewall Management Center.

You can apply a template to devices that are already registered with the Cloud-Delivered Firewall Management Center. Application of a template on a device clears existing configurations and applies configurations from the template. However, the Firewall Threat Defense HA failover configurations are not cleared.

Application of a template changes device configurations only on the Cloud-Delivered Firewall Management Center. You must explicitly deploy these device configuration changes to the Firewall Threat Defense device.

You cannot roll back the applied configuration changes. However, you can apply another template with the required configurations.

Note Any change management tickets related to template configurations must be approved for the corresponding changes to be incorporated into the template application workflow. Only approved template configurations are used during template application.

Follow these steps to apply a template to existing devices:

Procedure

Step 1 To apply the template from the **Template Management** window, choose **Devices** > **Template Management**.

a) Click the **More** (⋮) icon next to the template that you want to apply, and click **Apply**.
b) From the **Device** dropdown list, choose the **Device** on which you want to apply the template.
c) Click **Confirm** to initiate application of template on the device.

Step 2 (Optional) To apply the template from the **Associated Devices** window, choose **Devices > Template Management**.

a) Click the **Edit** (✎) icon of the template that you want to apply to a device.
b) Click **Associated Devices**.
c) In the **Associated Devices** window, click **Apply Template**.
d) From the **Device** drop-down list, choose the **Device** on which you want to apply the template.
e) Enter values for the **Variables** and **Network object overrides** fields.
f) Click **Apply** to initiate application of template on the device.

Reapply a template

Reapply a template to synchronize the device configuration with the template configuration when changes result in an out-of-sync state.

If you make any changes to the device or template that results in the configuration being out-of-sync, you can reapply the template to make the configuration in sync with the template.

Follow these steps to reapply a template on a device:

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template that you want to reapply to a device.

Step 3 Click **Associated Devices**.

Step 4 In the **Associated Devices** window, click **Reapply Template** for the device on which you want to reapply the template.

Note
If you want to reapply the template on all the associated devices in the template, click **Bulk Reapply** and click **Confirm**.

Step 5 On the **Reapply template** window, you can reuse the autopopulated **Variables** and **Network object overrides** values or enter new values.

Step 6 Click **Confirm** to initiate reapplication of the template on the device.

Monitor device templates

You can monitor and verify the application of templates by viewing the devices listed in the **Associated Devices** window and by viewing the **Template Apply Report**.

View associated devices

The devices that are associated with the templates are listed in the **Associated Devices** window. Each device row displays the **Device Name, Sync Status, Template Application Status**, and **Applied Date**. You can also click **Reapply template** to reapply the template. Click the **Variable Summary** icon to display a summary of the variables in the template and the **Report** ( ) icon to download the Device Template Apply report.

Click the **Delete** ( ) icon to remove the template from the device.

Click **Apply Template** if you want to apply a template on a device from the **Associated Devices** window. If you want to reapply the template on all the associated devices in the template, click **Bulk Reapply** and click **Confirm**.

The **Sync Status** can be either **Sync** or **Out-of-Sync**. **Sync** indicates that the template and device configurations are the identical. **Out-of-Sync** indicates that the configuration has changed on either the device or the template since the template was last applied.

These conditions do not affect the association of the device with the template:

- Pending configuration changes on the device—The **Sync Status** does not change if there are pending configuration changes that have to applied on the device.
- Deployment of pending configuration changes on the device—The **Sync Status** does not change after deployment of pending configuration changes on the device.

This table shows the **Sync** and **Out-of-Sync** scenarios that may occur.

Device Configurations Modified After Application of Template on Device	Template Configurations Modified After Application of Template on Device	Association Status
No	No	In Sync
Yes	No	Out of Sync
No	Yes	Out of Sync
Yes	Yes	Out of Sync

Template Apply reports

A **Template Apply Report** PDF is generated after the task to apply the template is completed. This report is generated on both successful and unsuccessful application of the template on the device. To access the link to this report, click **Notifications**, and then click the **Tasks** tab.

Template apply report details

The **Template Apply Report** contains these details:

- Template name
- Device model name
- Start and end time
- Status of the application of the template on the device
- Interface mapping information
- Variable values

Some configurations on the template may not be applied to the device if the device model or version is incompatible. The report includes details about these configurations and any errors encountered when the application of the template fails. Application of a template on a device may fail due to any of these reasons:

- Model mapping does not exist for the device model that is used.
- Values used for variables and network object overrides do not conform to routing policy or interface configuration rules. For example, the same IPv4 address has been used for two IPv4 address interface variables.
- The device or template is locked because another task, such as applying or modifying the template, is in progress.

Audit logs

View logs related to application of device templates, configuration updates, device template creation, and deletion activities that are recorded under audit logs.

Logs related to application of the device template, configuration updates, device template creation, and deletion, are logged under audit logs. The device template audit logs are added to the log both at the start and at the end of the task to apply the template on the device.

An audit diff file is also generated. This file contains the configuration changes that have been done during application of the template on the device.

Procedure

Step 1 Choose **System** (⚙) > **Monitoring** > **Audit**.

Step 2 The device template logs are logged under the subsystem **Devices** > **Template Management**. Click the **diff** icon to open a new window that displays the configuration changes that have been done during the application of the template on the device.

Troubleshoot device templates

This reference provides comprehensive troubleshooting information for device templates. For initial troubleshooting, we recommend looking at the information in the Template Apply Report and notifications that come up on the Cloud-Delivered Firewall Management Center UI when you run into an error. The Cloud-Delivered Firewall Management Center log files also contain detailed debugging and troubleshooting info.

Initial troubleshooting

Follow this procedure for initial troubleshooting.

1. Check the errors mentioned in the **Template Apply Report**. For more information, see Template Apply Report.
2. Review variable values and check for overlaps and incompatibilities.
3. Check model mappings to ensure if the correct model mappings exist. Delete or add mappings accordingly.
4. See the Cloud-Delivered Firewall Management Center audit logs to find any other issues and resolve them.

Consider this error scenario. In a device template, the inside interface is configured with a static IPv4 variable - *$insideIPv4*.

The BGP IPv4 address is configured with an IPv4 BGP neighbor.

An overlapping IPv4 address is configured for the BGP neighbor and an interface.

Due to these issues, the application of the device template fails and an error is displayed.

To troubleshoot this error, identify the error from the notification displayed on the UI.

```
IP Address 192.168.10.1 same as ip address of interface - 'inside'(Ethernet1/1)
```

Check the **Template Apply Report** for more information.

Enter correct values for the variables and apply the template again to ensure successful application of the template on the device.

Troubleshoot device registration

- Issue: Admin Password is incorrect or not provided during registration

  Scenario: If the admin password is not set on the device and if you have not provided the admin password during registration, the Firewall Threat Defense device provisioning will fail. In such a scenario, a *Provision Error* along with an **Enter Password** link is displayed.

  Workaround: Click **Enter Password** to enter a new password and click **Save**. Click **Confirm and Proceed** to trigger the onboarding again.

- If the admin password is already set on the device and you provide another admin password during registration, device provisioning will fail.
- Issue: Device registration in Cloud-Delivered Firewall Management Center fails

  Workaround: Follow existing device registration troubleshooting steps. For more information, see Configure, Verify, and Troubleshoot Firepower Device Registration.
- Issue: Bulk Registration Request Fails in Cloud-Delivered Firewall Management Center

  Scenario: The bulk registration request can fail due to a few scenarios:
  - You do not have the required permissions to perform template-related operations
  - Invalid CSV file provided

  Workaround: You can see logs for these errors in the VMS Shared and USM Shared log files. Fix the errors and initiate registration again.
- Issue: Device provisioning fails in Security Cloud Control due to some generic errors, such as communication with the device fails

  Workaround: Click Retry in the Provision Error to trigger the onboarding in Security Cloud Control again. You can also see the Security Cloud Control workflows for more information on the error and troubleshooting information.

Troubleshoot Cisco Security Cloud integration

Issue: Cisco Security cloud integration not successful

Workaround: Follow Cisco Security Cloud integration troubleshooting steps. For more information, see Cisco Security Cloud Integration.

Troubleshoot device template configuration issues

Issue: Device template misconfigurations causing deployment failures after registration

Workaround: Follow these steps for initial troubleshooting.

1. Check the errors mentioned in the Template Apply Report.
2. Review variable values and check for overlaps and incompatibilities.
3. Check model mappings to ensure if the correct model mappings exist. Delete or add mappings accordingly.
4. See the Cloud-Delivered Firewall Management Center audit logs to find any other issues and resolve them.

Troubleshoot Security Cloud Control Issues

- Issue: Device with serial number already claimed

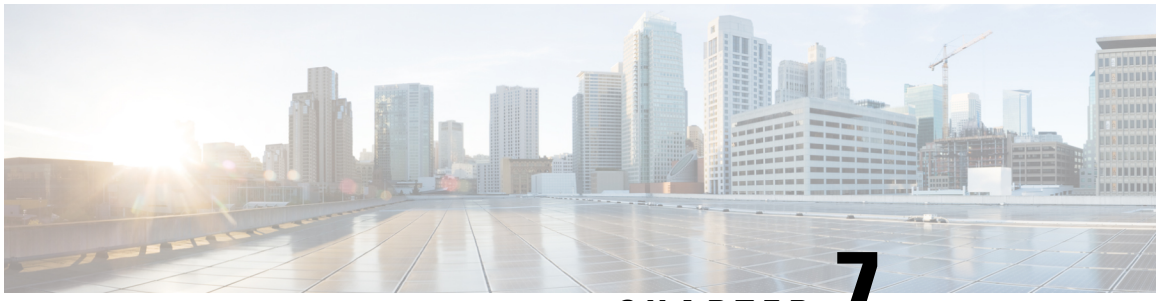
  Workaround: Verify serial number and reinitiate onboarding.
- Issue: Security Cloud Control fails to claim devices

  Workaround: Select the device in the Security Cloud Control **Security Devices** window for more details on the error. You can see logs related to device claim issues in the VMS Shared and USM Shared log files. Click **Retry** to initiate registration again.

- Issue: Communication failures between Cloud-Delivered Firewall Management Center and Security Cloud Control

  Scenario: Communication failures between Cloud-Delivered Firewall Management Center and Security Cloud Control can cause failures during the Zero-Touch Provisioning (ZTP) device registration request.

  Workaround: Refresh the ZTP device status, retry ZTP registration, and delete the ZTP device. You can see logs regarding communication failure between the Cloud-Delivered Firewall Management Center and Security Cloud Control in the Auth Daemon logs. For operational failures related to ZTP, you can see the logs in the VMS Shared and USM Shared log files.



CHAPTER 7

Transparent or Routed Firewall Mode

This chapter describes how to set the firewall mode to routed or transparent, as well as how the firewall works in each firewall mode.

Note The firewall mode only affects regular firewall interfaces, and not IPS-only interfaces such as inline sets or passive interfaces. IPS-only interfaces can be used in both firewall modes. See Inline Sets and Passive Interfaces, on page 1119 for more information about IPS-only interfaces. Inline sets might be familiar to you as "transparent inline sets," but the inline interface type is unrelated to the transparent firewall mode described in this chapter or the firewall-type interfaces.

Caution

- Use FTD CLI commands to set "Firewall Mode."

- Firewall mode, on page 153
- Default settings, on page 163
- Guidelines for firewall mode, on page 163
- Set the firewall mode, on page 164

Firewall mode

A Firewall Threat Defense device supports two firewall modes for regular firewall interfaces:

- Routed Firewall mode
- Transparent Firewall mode

Routed firewall mode

In routed mode, the Firewall Threat Defense device is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet.

Integrated routing and bridging features

With Integrated Routing and Bridging, you can use a "bridge group" where you group together multiple interfaces on a network, and the Firewall Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. The Firewall Threat Defense device routes between BVIs and regular routed interfaces. If you do not need clustering or EtherChannel member interfaces, you might consider using routed mode instead of transparent mode. In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces as well for a mixed deployment.

Transparent firewall mode

Traditionally, a firewall is a routed hop and acts as a default gateway for hosts that connect to one of its screened subnets. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices. However, like any other firewall, access control between interfaces is controlled, and all of the usual firewall checks are in place.

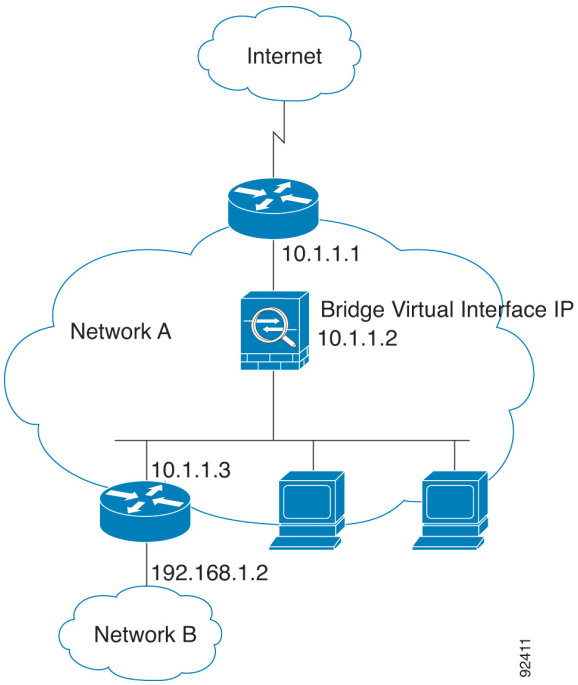
Layer 2 connectivity is achieved by using a "bridge group" where you group together the inside and outside interfaces for a network, and the Firewall Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. You can have multiple bridge groups for multiple networks. In transparent mode, these bridge groups cannot communicate with each other.

Transparent firewall network usage

The Firewall Threat Defense device connects the same network between its interfaces. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network.

The figure shows a typical transparent firewall network where the outside devices are on the same subnet as the inside devices. The inside router and hosts appear to be directly connected to the outside router.

Figure 9: Transparent firewall network



Passing traffic for routed-mode features

For features that are not directly supported on the transparent firewall, you can allow traffic to pass through so that upstream and downstream routers can support the functionality.

Supported traffic types

For example, by using an access rule, you can allow DHCP traffic (instead of the unsupported DHCP relay feature) or multicast traffic such as that created by IP/TV. You can also establish routing protocol adjacencies through a transparent firewall; you can allow OSPF, RIP, EIGRP, or BGP traffic through based on an access rule. Likewise, protocols like HSRP or VRRP can pass through the Firewall Threat Defense device.

Bridge groups

A bridge group is a group of interfaces that the Firewall Threat Defense device bridges instead of routes. Bridge groups are supported in both transparent and routed firewall mode. Like any other firewall interfaces, access control between interfaces is controlled, and all of the usual firewall checks are in place.

Bridge virtual interfaces

Each bridge group includes a Bridge Virtual Interface (BVI). The Firewall Threat Defense device uses the BVI IP address as the source address for packets originating from the bridge group. The BVI IP address must be on the same subnet as the bridge group member interfaces. The BVI does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.

BVI functionality by mode

In transparent mode: Only bridge group member interfaces are named and can be used with interface-based features.

In routed mode: The BVI acts as the gateway between the bridge group and other routed interfaces. To route between bridge groups/routed interfaces, you must name the BVI. For some interface-based features, you can use the BVI itself:

- DHCPv4 server—Only the BVI supports the DHCPv4 server configuration.
- Static routes—You can configure static routes for the BVI; you cannot configure static routes for the member interfaces.
- Syslog server and other traffic sourced from the Firewall Threat Defense device—When specifying a syslog server (or SNMP server, or other service where the traffic is sourced from the Firewall Threat Defense device), you can specify either the BVI or a member interface.

If you do not name the BVI in routed mode, then the Firewall Threat Defense device does not route bridge group traffic. This configuration replicates transparent firewall mode for the bridge group. If you do not need clustering or EtherChannel member interfaces, you might consider using routed mode instead. In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces as well for a mixed deployment.

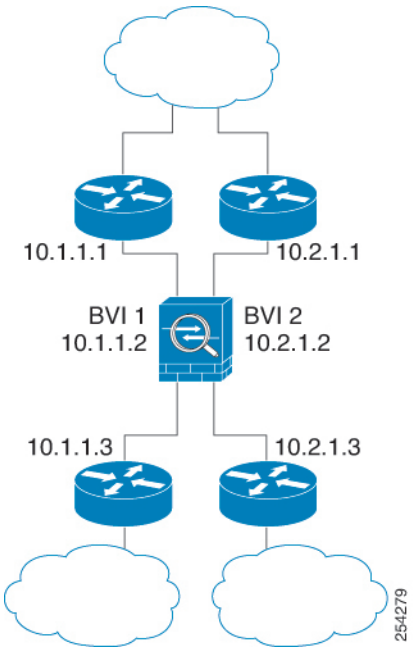
Bridge groups in transparent firewall mode

Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the Firewall Threat Defense device , and traffic must exit the Firewall Threat Defense device before it is routed by an external router back to another bridge group in the Firewall Threat Defense device. Although the bridging functions are separate for each bridge group, many other functions are shared between all bridge groups. For example, all bridge groups share a syslog server or AAA server configuration.

You can include multiple interfaces per bridge group. See Guidelines for firewall mode, on page 163 for the exact number of bridge groups and interfaces supported. If you use more than 2 interfaces per bridge group, you can control communication between multiple segments on the same network, and not just between inside and outside. For example, if you have three inside segments that you do not want to communicate with each other, you can put each segment on a separate interface, and only allow them to communicate with the outside interface. Or you can customize the access rules between interfaces to allow only as much access as desired.

This figure shows two networks connected to the Firewall Threat Defense device, which has two bridge groups.

Figure 10: Transparent firewall network with two bridge groups



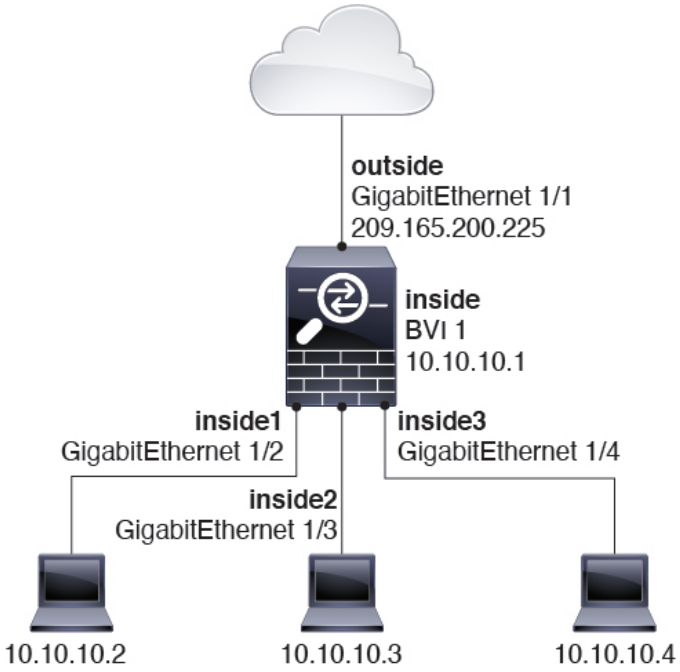
Bridge groups in routed firewall mode

Bridge group traffic can be routed to other bridge groups or routed interfaces. You can choose to isolate bridge group traffic by not assigning a name to the BVI interface for the bridge group. If you name the BVI, then the BVI participates in routing like any other regular interface.

Bridge group usage and configuration

One use for a bridge group in routed mode is to use extra interfaces on the Firewall Threat Defense instead of an external switch. For example, the default configuration for some devices include an outside interface as a regular interface, and then all other interfaces assigned to the inside bridge group. Because the purpose of this bridge group is to replace an external switch, you need to configure an access policy so all bridge group interfaces can freely communicate.

Figure 11: Routed firewall network with an inside bridge group and an outside routed interface



Allow layer 3 traffic

This topic describes how different types of Layer 3 traffic are handled through bridge groups and the access rule requirements for each traffic type.

- Unicast IPv4 and IPv6 traffic requires an access rule to be allowed through the bridge group.
- ARPs are allowed through the bridge group in both directions without an access rule. ARP traffic can be controlled by ARP inspection.
- IPv6 neighbor discovery and router solicitation packets can be passed using access rules.
- Broadcast and multicast traffic can be passed using access rules.

Allowed MAC addresses

These destination MAC addresses are allowed through the bridge group if allowed by your access policy (For more information, refer to Allow layer 3 traffic, on page 158). Any MAC address not on this list is dropped.

- TRUE broadcast destination MAC address equal to FFFF.FFFF.FFFF
- IPv4 multicast MAC addresses from 0100.5E00.0000 to 0100.5EFE.FFFF
- IPv6 multicast MAC addresses from 3333.0000.0000 to 3333.FFFF.FFFF
- BPDU multicast address equal to 0100.0CCC.CCCD

BPDUs handling

To prevent loops using the Spanning Tree Protocol, BPDUs are passed by default.

By default BPDUs are also forwarded for advanced inspection, which is unnecessary for this type of packet, and which can cause problems if they are blocked due to an inspection restart, for example. We recommend that you always exempt BPDUs from advanced inspection. To do so, use FlexConfig to configure an EtherType ACL that trusts BPDUs and exempts them from advanced inspection on each member interface.

The FlexConfig object should deploy the following commands, where you replace <if-name> with an interface name. Add as many access-group commands as needed to cover each bridge group member interface on the device. You can also choose a different name for the ACL.

```
access-list permit-bpdu ethertype trust bpdu
access-group permit-bpdu in interface <if-name>
```

MAC address and route lookups

For traffic within a bridge group, the outgoing interface of a packet is determined by performing a destination MAC address lookup instead of a route lookup.

Route lookups, however, are necessary for these situations:

- Traffic originating on the Firewall Threat Defense device—Add a default/static route on the Firewall Threat Defense device for traffic destined for a remote network where a syslog server, for example, is located.
- SSL policy enabled in transparent firewall mode—
  - TLS 1.3 traffic passes through the Firewall Threat Defense device device.
  - Firewall Threat Defense device acts as a TCP proxy. It terminates the client-side TCP connection and initiates a separate server-side connection. The Firewall Threat Defense device TCP stack owns the session state.
  - For idle proxied sessions, Firewall Threat Defense device generates TCP ACK keepalive probes to preserve the connection state. These are not passthrough packets—the firewall originates them and sends them toward the client or server depending on the state of the TLS transit connection.
  - Because the Firewall Threat Defense device originates these keepalives, it must perform an adjacency table lookup (next-hop or egress MAC resolution) even in transparent mode.
  - If the adjacency lookup fails for one of these proxy-generated packets, the Firewall Threat Defense device logs syslog 110003.

    To prevent these syslog messages, configure routes on the Firewall Threat Defense device for the transit connection endpoints (client and server IP addresses). Proper routing enables successful adjacency lookups for keepalive packets.

    If the Firewall Threat Defense device successfully sends a TCP keepalive and receives a reply, it keeps the transit connection active. If the device cannot send the keepalives, or receives no reply, it removes the transit connection after approximately 240 seconds.
- Voice over IP (VoIP) and TFTP traffic, and the endpoint is at least one hop away—Add a static route on the Firewall Threat Defense device for traffic destined for the remote endpoint so that secondary connections are successful. The Firewall Threat Defense device creates a temporary "pinhole" in the access control policy to allow the secondary connection; and because the connection might use a different set of IP addresses than the primary connection, the Firewall Threat Defense device needs to perform a route lookup to install the pinhole on the correct interface.

Affected applications include:

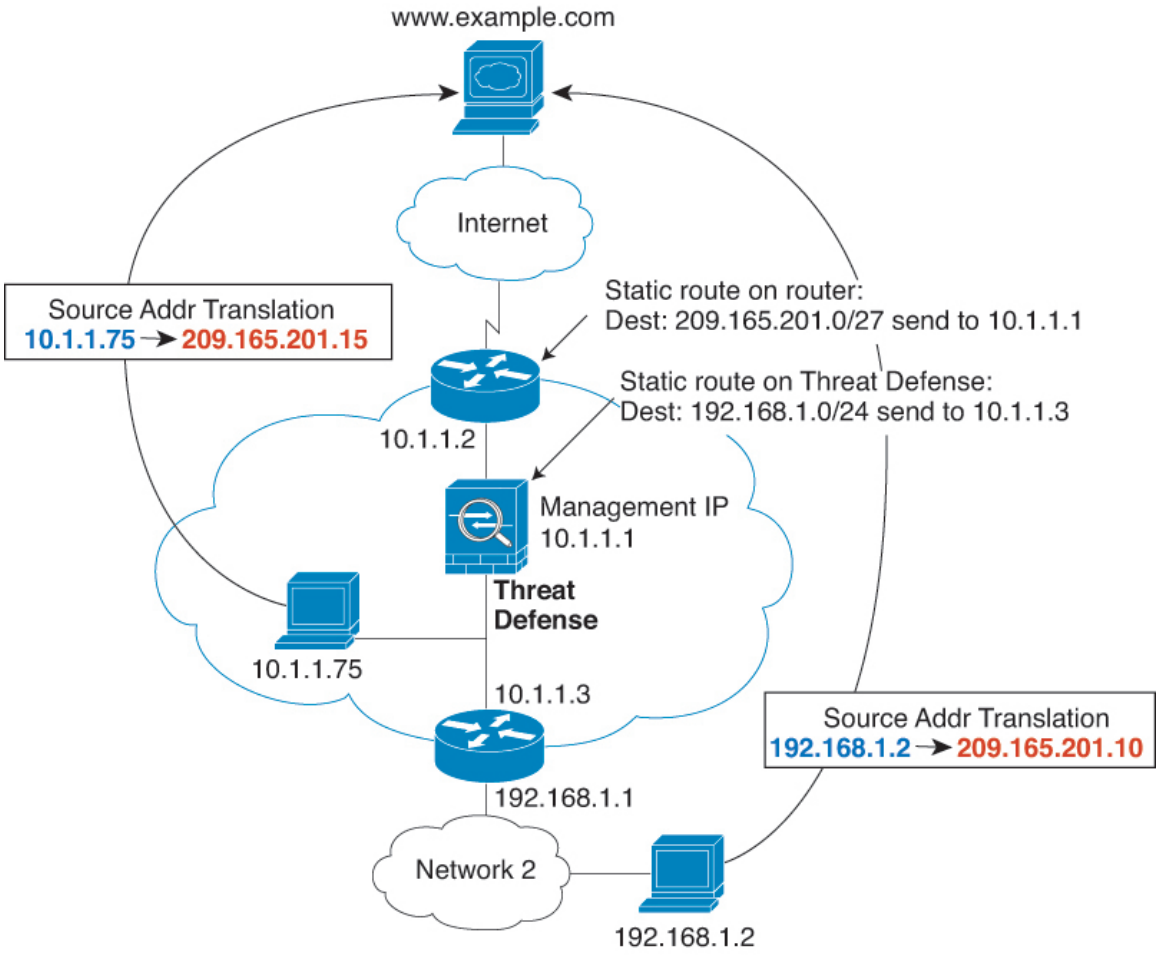
- H.323
- RTSP
- SIP and Skinny (SCCP)
- SQL*Net
- SunRPC
- TFTP

- Traffic at least one hop away for which the Firewall Threat Defense device performs NAT—Configure a static route on the Firewall Threat Defense device for traffic destined for the remote network. You also need a static route on the upstream router for traffic destined for the mapped addresses to be sent to the Firewall Threat Defense device.

  This routing requirement is also true for embedded IP addresses for VoIP and DNS with NAT enabled, and the embedded IP addresses are at least one hop away. The Firewall Threat Defense device needs to identify the correct egress interface so it can perform the translation.

NAT within a bridge group

Figure 12: NAT example: NAT within a bridge group



Unsupported features for bridge groups in transparent mode

This table lists the features that are not supported in bridge groups in transparent mode.

Table 4: Unsupported features in transparent mode

Feature	Description
Dynamic DNS	—
DHCPv6 stateless server	Only the DHCPv4 server is supported on bridge group member interfaces.

Feature	Description
DHCP relay	The transparent firewall can act as a DHCPv4 server, but it does not support DHCP relay. DHCP relay is not required because you can allow DHCP traffic to pass through using two access rules: one that allows DCHP requests from the inside interface to the outside, and one that allows the replies from the server in the other direction.
Dynamic routing protocols	You can, however, add static routes for traffic originating on the Firewall Threat Defense device for bridge group member interfaces. You can also allow dynamic routing protocols through the Firewall Threat Defense device using an access rule.
Multicast IP routing	You can allow multicast traffic through the Firewall Threat Defense device by allowing it in an access rule.
QoS	—
VPN termination for through traffic	The transparent firewall supports site-to-site VPN tunnels for management connections only on bridge group member interfaces. It does not terminate VPN connections for traffic through the Firewall Threat Defense device. You can pass VPN traffic through the Firewall Threat Defense device using an access rule, but it does not terminate non-management connections.

Unsupported features for bridge groups in routed mode

This table lists the features are not supported in bridge groups in routed mode.

Table 5: Unsupported features in routed mode

Feature	Description
EtherChannel member interfaces	Only physical interfaces, redundant interfaces, and subinterfaces are supported as bridge group member interfaces. Management interfaces are also not supported.
Clustering	Bridge groups are not supported in clustering.
Dynamic DNS	—
DHCPv6 stateless server	Only the DHCPv4 server is supported on BVIs.
DHCP relay	The routed firewall can act as a DHCPv4 server, but it does not support DHCP relay on BVIs or bridge group member interfaces.
Dynamic routing protocols	You can, however, add static routes for BVIs. You can also allow dynamic routing protocols through the Firewall Threat Defense device using an access rule. Non-bridge group interfaces support dynamic routing.
Multicast IP routing	You can allow multicast traffic through the Firewall Threat Defense device by allowing it in an access rule. Non-bridge group interfaces support multicast routing.

Feature	Description
QoS	Non-bridge group interfaces support QoS.
VPN termination for through traffic	You cannot terminate a VPN connection on the BVI. Non-bridge group interfaces support VPN. Bridge group member interfaces support site-to-site VPN tunnels for management connections only. It does not terminate VPN connections for traffic through the Firewall Threat Defense device. You can pass VPN traffic through the bridge group using an access rule, but it does not terminate non-management connections.

Default settings

Lists the default settings for firewall mode.

Bridge group defaults

By default, all ARP packets are passed within the bridge group.

Guidelines for firewall mode

Review these guidelines for firewall mode.

Bridge group limits and network requirements

- You can create up to 250 bridge groups, with 64 interfaces per bridge group.
- Each directly-connected network must be on the same subnet.
- The Firewall Threat Defense device does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.

BVI IP address requirements

- An IP address for the BVI is required for each bridge group for to-the-device and from-the-device management traffic, as well as for data traffic to pass through the Firewall Threat Defense device. For IPv4 traffic, specify an IPv4 address. For IPv6 traffic, specify an IPv6 address.
- You can only configure IPv6 addresses manually.
- The BVI IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255).

Unsupported bridge group members

- Management interfaces are not supported as bridge group members.

- For multi-instance mode, shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode).
- For the Firewall Threat Defense Virtual on VMware with bridged ixgbevf interfaces, transparent mode is not supported, and bridge groups are not supported in routed mode.
-
- For the 200/ 1010/ 1210// 1220, you cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.
- For the Firepower 4100/9300, data-sharing interfaces are not supported as bridge group members.

Transparent mode requirements and limitations

- In transparent mode, you must use at least 1 bridge group; data interfaces must belong to a bridge group.
- In transparent mode, do not specify the BVI IP address as the default gateway for connected devices; devices need to specify the router on the other side of the Firewall Threat Defense as the default gateway.
- In transparent mode, the default route, which is required to provide a return path for management traffic, is only applied to management traffic from one bridge group network. This is because the default route specifies an interface in the bridge group as well as the router IP address on the bridge group network, and you can only define one default route. If you have management traffic from more than one bridge group network, you need to specify a regular static route that identifies the network from which you expect management traffic.
- Transparent mode is not supported on threat defense virtual instances deployed on Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud Infrastructure.

Routed mode requirements and limitations

- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, Firewall Threat Defense-defined EtherChannel interfaces are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.

Traffic handling limitation

- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the Firewall Threat Defense when using bridge group members. If there are two neighbors on either side of the Firewall Threat Defense running BFD, then the Firewall Threat Defense will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

Set the firewall mode

You can set the firewall mode when you perform the initial system setup at the CLI. We recommend setting the firewall mode during setup because changing the firewall mode erases your configuration to ensure you do not have incompatible settings. If you need to change the firewall mode later, you must do so from the CLI.

Procedure

Step 1 Unregister the Firewall Threat Defense device from the Cloud-Delivered Firewall Management Center.

You cannot change the mode until you unregister the device.

a) Performance profiles are retained across firewall mode changes and can't be changed after you re-register. You must set the performance profile to Default. See the **Performance Profile** page in your platform settings policy (choose **Devices** > **Platform Settings**, and then click **Edit** (✎) for your policy).

Deploy the configuration to make these changes.

b) Choose **Devices** > **Device Management**.

c) Next to the device you want to unregister, click **More** (⋮), and then click **Delete**.

d) If you have switched to the new device management page, select the device to open the side pane, and under **Device actions**, click **Delete**.

Step 2 Access the Firewall Threat Defense device CLI, preferably from the console port.

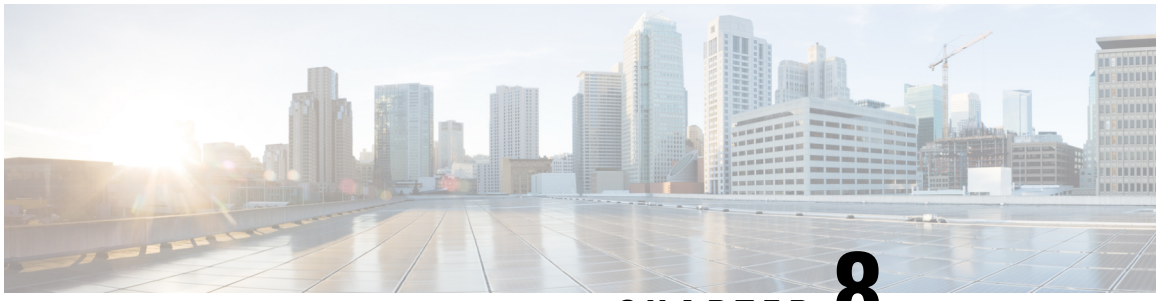
Step 3 Change the firewall mode:

configure firewall [**routed** | **transparent**]

Example:

```
> configure firewall transparent
This will destroy the current interface configurations, are you sure that you want to
proceed? [y/N] y
The firewall mode was changed successfully.
```

Step 4 Re-register with the Cloud-Delivered Firewall Management Center.



CHAPTER 8

Change Management

You can enable change management if your organization needs to implement more formal processes for configuration changes, including audit tracking and official approval before changes are deployed.

- Change management, on page 167
- Requirements and prerequisites for change management, on page 171
- Guidelines and limitations for change management, on page 172
- Enabling or disabling change management, on page 172
- Manage tickets, on page 173

Change management

Change management is a formal process that

- enforces an official approval process before configuration changes can be made to a device,
- requires administrators to open tickets before they can make configuration changes, and
- ensures that the right employees make the final decisions through submission and approval of tickets.

Change management characteristics

When using change management, the system implements these controls:

- Administrators can see their own changes within a ticket, but they cannot see changes anyone else has made within a ticket
- A policy is locked once a user makes a change within a ticket, preventing interfering changes
- Users cannot make changes while another user has made a change that is pending approval

Administrators can create multiple tickets so that a single ticket contains only logically-related policy changes. Tickets with a more limited scope are also easier to evaluate and approve quickly.

Make configuration changes in the change management workflow

This task enables configuration specialists to make configuration changes to supported policies and objects when change management is enabled.

When you enable change management, users who configure devices and shared policies need to change their approach slightly. Configuration specialists need to take the following approach when making configuration changes to supported policies and objects.

Procedure

Step 1 Create a ticket.

Step 2 Open the ticket.

Step 3 Make the configuration changes.

Note that the procedures explained in the online help and user guides assume that change management is not active, and omit any steps for creating, opening, or submitting tickets.

Step 4 Optionally, preview and validate the ticket to ensure the changes are complete and correct.

Step 5 Submit the ticket. At this point, the approver can either approve or reject the ticket.

- If the ticket is approved, deploy the changes.
- If the ticket is rejected, address the issues, and resubmit the ticket.

The configuration changes are successfully implemented through the change management workflow with proper approval and deployment processes.

Create separate approver and configuration roles

Separate approver and configuration roles are system-defined and custom roles that

- have permissions to modify (create/open/discard) and review (approve/reject/take over) tickets,
- provide granular access control to separate ticket approval from configuration changes, and
- ensure that ticket approval is assigned only to users who have organizational authority to approve changes.

System-defined roles and permissions

Some system-defined roles have permissions to modify (create/open/discard) and review (approve/reject) tickets:

- To both modify and review tickets:
  - Admin
- To modify tickets only:
  - Edit Only
- To review tickets only:
  - Deploy Only

The **Deploy Only** role can approve or reject change management tickets. Approving a ticket makes the ticketed changes available for deployment, but does not deploy the changes to the device. You must deploy approved changes separately.

Note A Read Only user cannot use the change management feature.

If you need more granular roles to separate these activities due to your organizational requirements, you can create separate roles to ensure that ticket approval is assigned only to those users who have the organizational authority to approve changes. To create a new user, navigate back to Security Cloud Control and from the Security Cloud Control navigation bar, choose **Settings** > **User Management**.

The approach you take depends on your precise requirements. For example:

- If your approvers should also be allowed to make configuration changes, you can simply assign them the system-defined roles, such as Administrator. Then, create custom configuration-only roles that include the same permissions but not the Review Tickets permission.
- If you need complete separation between approvers and those who make configuration changes, create custom roles for both, limiting the roles to either the Modify Tickets or the Review Tickets permission plus all other needed permissions for viewing or changing the supported policies and objects.

Policies and objects that support change management

If a policy or object supports the change management workflow, then creating, editing, or deleting the policy or object, including assigning a policy to a device, must be done in an open ticket.

Any action, policy, or object that does not support the change management workflow can be created, edited, or deleted, and so forth, without an open ticket. Even if a ticket is open, the changes made to unsupported policies are not included in the ticketed changes and are available for deployment immediately.

This information includes the policies and objects that are supported. Anything not listed is unsupported.

Supported policies

- Access control, including rules, references to other policies, and inheritance settings. Cloning access control policies is not included within a ticket. The clone is immediately accepted and available to all users.
- Device configuration policies:
  - Interfaces
  - Inline sets
  - DHCP
  - VTEP
  - All routing
- Decryption policy

- DNS policy
- FlexConfig
- Intrusion policy and network analysis policy (NAP), Snort 3 only.
- Malware and file policy
- Network address translation (NAT)
- Network discovery policy
- Platform settings
- Prefilter policy
- QoS
- Splunk integration
- Umbrella SASE topology
- VPN policies, both site-to-site and remote access
- Zero trust access

Supported objects

- AAA server
- Access list
- Address pools
- AS path
- Cipher suit lists
- Community list
- Distinguished name objects
- DHCP IPv6 pools
- DNS server group
- FlexConfig objects
- Group policy
- Interface
- Key chain
- Network
- PKI certificates, all objects
- Policy list
- Port

- Prefix list
- PKI certificates, all objects
- Route map
- Sinkhole
- SLA monitor
- Time range
- Time zone
- Tunnel zone
- URL
- Variable set
- VLAN tag
- VPN objects (IKEv1, IKEv2 IPSec and policy, PKI enrollment, certificate map)

Requirements and prerequisites for change management

This reference provides the requirements and prerequisites necessary to configure change management functionality.

Model support

Management center

User roles

- To enable or disable change management: Admin.
- To both modify and review tickets:
  - Admin
- To modify tickets only:
  - Edit Only
- To review tickets only:
  - Deploy Only

    The **Deploy Only** role can approve or reject change management tickets. Approving a ticket makes the ticketed changes available for deployment, but does not deploy the changes to the device. You must deploy approved changes separately.

Guidelines and limitations for change management

When operating in change management mode, users can make changes to supported policies, but they cannot save the changes. For example, you could go through the dialog box to create a new platform settings policy without an open ticket, but when you click OK to actually create the policy, you will get an error and the policy will not be created.

The following activities require that all tickets be in a terminal state, that is, approved or discarded: backup/restore.

Deleting a device from the inventory requires that all tickets involving that device be approved or discarded.

Some system processes prevent you from enabling or disabling change management. Wait until the following processes complete before you enable or disable change management:

- critical management center processes, including SRU updates, LSP updates, VDB updates, and management center upgrades
- pending or running background activities, including device registration; cluster auto-registration; Firewall Threat Defense high availability operations; high availability auto-registration; FlexConfig migration; cluster creation, break, edit, reconcile, or node changes; device-template creation or application; device-model migration; and EPM join or breakout
- system jobs, including backup or restore, device import or export, general import, and domain movement.

Your ability to create objects while configuring a feature is constrained based on whether the feature and objects are all supported by change management. For example, importing a configuration is not supported by change management. Therefore, you cannot create security zone objects, which are supported, during the import. On the other hand, you can create new objects while configuring access control rules, because both are supported.

When you create an object from Security Cloud Control, the system automatically creates a ticket internally and allows the object to be associated with the Cloud-Delivered Firewall Management Center. You do not have to create or have an open ticket to do this. However, when you want to create an object from the Cloud-Delivered Firewall Management Center, you need an existing ticket or create one. The object is synchronized to Security Cloud Control only after the ticket is approved.

When using Cloud-Delivered Firewall Management Center, a user defined in Security Cloud Control is available to be assigned tickets only after the user cross-launches cdFMC at least once. Until the first cross-launch, the user does not exist in cdFMC.

Enabling or disabling change management

Enable or disable change management to control whether users must open tickets and get approval when making configuration changes to the system.

The default is that the change management workflow is disabled. Users do not need to open tickets and get approval when making configuration changes. If you want to enforce the change management workflow, you must enable it globally for the system.

Before you begin

There are several system processes that prevent you from enabling/disabling change management. If any of the following are in process, you need to wait for them to complete before changing these settings: backup/restore; import/export; domain movement; upgrade; Flexconfig migration; device registration; high-availability registration, creation, break, or switch; cluster create, registration, break, edit, add or remove nodes; EPM break out or join.

An access control policy cannot be locked when you change these settings. If a policy is locked, you must wait for the lock to be released before enabling/disabling this feature.

Follow these steps to enable or disable change management:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Change Management**.

Step 3 Select **Enable Change Management**.

To disable the feature, deselect the option. All tickets must be approved or discarded to disable change management. You cannot disable change management if any ticket is in the In Progress, On Hold, Rejected, or Pending Approval state.

Step 4 Select the **Number of approvals required**, which is how many administrators must approve the change for the ticket to be approved and deployable. The default is 1, but you can require up to 5 approvers per ticket. Users can override this number when creating tickets.

Step 5 Select the **Ticket Purge Duration**, which is the number of days to keep approved tickets, from 1-100 days. The default is 5 days.

Step 6 (Optional.) Enter the **Reply to Address** and the email addresses for the **List of Approver Addresses**. You must also configure the email notification system settings for email to work.

Step 7 Click **Save**.

The system adds the **Ticket** (📄) shortcut to the menu bar, and the **System** (⚙) > **Change Management Workflow** command. Users can manage tickets using these methods.

Change management is enabled or disabled according to your configuration. If enabled, the system adds ticket management shortcuts to the interface, and users must follow the change management workflow when making configuration changes.

Manage tickets

This task enables you to manage change management tickets for configuration changes to supported policies that require approval workflow.

When you enable change management, configuration changes for supported policies must be done within the context of a ticket. You open a ticket, make your changes, then submit the ticket for approval.

You can see a list of tickets, and create new ones, either on the change management page or through the ticket quick-access menu. All ticket changes are synchronized in each menu, so you can switch back and forth at your convenience and use whichever method you prefer.

Note When you open a ticket and make a change to a supported policy, that policy is locked from changes by other users or through other tickets. The policy remains locked until the ticket is approved or discarded.

Procedure

Step 1 Do any of these actions:

- Choose **System** (⚙) > **Change Management Workflow** to open a page showing existing tickets.
- Click the **Ticket** (📋) quick-access menu. The icon can either be named Select a Ticket (if no ticket is opened), the ticket name if a ticket is opened, or unnamed if no tickets exist.

Both pages are organized the same. The **Ticket** tab lists all tickets, whereas the **Review** tab lists tickets that have been submitted for approval. The default view shows your tickets only.

Step 2 On the **Ticket** tab, take any of these actions:

- To create a new ticket, click **Add Ticket**.
- To view the details of a ticket, click the > next to the ticket name. The Details page includes UUID, name, description, user, last modified date, and comments. The History page includes the status changes for the ticket. The image at the top shows where the ticket stands in the overall workflow.
- To preview the configuration changes for an open ticket, click **Preview** (📄).
- To validate the configuration changes in an open ticket, click **Validate** (📋) or **More** (⋮) > **Validate**. A dialog box opens showing error, warning, and informational messages, if there are any validation errors.
- To open a ticket, click **Open** (📂) or **More** (⋮) > **Open**. To close an open ticket, click **Put Ticket on Hold (X)** or **More** (⋮) > **Put Ticket on Hold**. Closing a ticket does not submit it for review, nor does it release any locks placed on edited policies.
- To submit an open ticket for review and approval, click **Submit for Approval** (⬆) or **More** (⋮) > **Submit for Approval**. The ticket must be open to be submitted.
- To discard a ticket, click **Discard** (🗑) or **More** (⋮) > **Discard**.
- To take over or reassign a ticket, click **More** (⋮) > **Take Over Ticket** when viewing all tickets in the system.
- To search for a ticket, type a string in the search box. The search looks at ticket name, description, and responsible user. To filter the list by ticket status (on the change management workflow page), click the status above the list: **New**, **Open**, **On Hold** (ticket was closed), **Rejected**, **Pending Approval**, **Approved**. Each status has a count of the number of tickets in that state. Click **All** under **My Tickets** to return to the default of showing all of your tickets, or **All** under **Tickets in System** to see everyone's tickets.

Step 3 On the **Reviews** tab, take any of these actions on submitted tickets.

The list is empty if there are no submitted tickets. In addition, only users with Review Ticket permissions can see this tab.

- To preview the configuration changes for a ticket, click **Preview** ( ).
- To validate the configuration changes in an open ticket, click **Validate** ( ) or **More** ( ) > **Validate**.
- To approve the ticket, click **Approve** ( ) or **More** ( ) > **Approve**.
- To disapprove the ticket, click **Reject** ( ) or **More** ( ) > **Reject**.

You have successfully managed your change management tickets. Approved tickets can proceed with configuration changes, while rejected tickets require revision before resubmission.

Create change management tickets

Create change management tickets to ensure all configuration changes are performed within the context of an approved workflow ticket.

When using the change management workflow, you must do all configuration changes within the context of an open ticket. If you do not already have a ticket, you must create a new one.

Procedure

Step 1 Choose **System** ( ) > **Change Management Workflow**, or click the **Ticket** ( ) shortcut menu.

Step 2 Click **Add Ticket**.

Step 3 Configure the ticket options:

- **Name**—The name of the ticket. The name can include letters, numbers, spaces, and the following special characters: #-_!
- **Description**—An optional description of what you intend to configure using this ticket. For example, if you have a case number related to what you intend to fix using this ticket, that would be useful information in a description.
- **Number of Approvers**—How many administrators must approve the change for the ticket to be approved and deployable. You can specify 1-5.
- **Assign to**—Select the user who will own the ticket and be responsible for implementing the changes. Choose **self** to assign it to yourself.

Step 4 Click one of the following:

- **Create**—The ticket is added to the list of tickets, but it is not opened. You need to open it before you can work within the context of the ticket.
- **Create and Open**—The ticket is added to the list of tickets and also opened.

The change management ticket is created and available for use in the workflow. If you selected Create and Open, the ticket is immediately opened for configuration changes.

Opening a ticket for configuration changes

Open a ticket to enable configuration changes within the ticket management system.

Before you can make changes within a ticket, you must open the ticket.

If you have another ticket open, the system puts it on hold (closes it) for you before opening the new one.

Procedure

Step 1 Choose **System** ( ) > **Change Management Workflow**, or click the **Ticket** ( ) shortcut menu.

Step 2 On the **Tickets** tab, click **Open** ( ) or **More** ( ) > **Open** for the ticket.

Step 3 Optionally, enter a comment for the action.

Step 4 Click **Open**.

You can now start your configuration changes. The name of the ticket icon changes to the name of the open ticket.

The ticket is now open and ready for configuration changes. The ticket icon displays the name of the open ticket.

Preview a ticket

Preview a ticket to view all configuration changes made within the context of the ticket.

You can preview a ticket while making your configuration changes, or before approving it. The preview shows all configuration changes made within the context of the ticket.

Procedure

Step 1 Choose **System** ( ) > **Change Management Workflow**, or click the **Ticket** ( ) shortcut menu.

Step 2 Click **Preview** ( ) for the ticket.

The Preview dialog box opens. Changes are color coded according to the legend at the top of the dialog box.

Step 3 Select the policy whose changes you want to view in the Changed Policies list.

You are shown both the current version of the policy as it is defined in Cloud-Delivered Firewall Management Center (on the left) and the proposed changes defined within the ticket.

For policies that contain pages, such as platform settings, you can select the overall policy to see all changes, or specific pages within the policy, in the Changed Policies list.

You cannot alter the changes from within the preview. If you need to change something, you must close the preview and return to the policy you want to change.

Step 4 Optionally, click **Download as PDF** to save the preview to a PDF file for offline viewing or archival purposes.

Step 5 Click **OK**.

The ticket preview displays all configuration changes, allowing you to review them before approval or while making modifications.

Submit a ticket

Submit a ticket to apply changes made within the ticket after validation and review.

After you have completed the changes needed for a ticket, you can preview and validate the changes. Then, when you are satisfied with the changes, submit the ticket for review and approval.

Changes made within the ticket are not applied until you submit the ticket and the ticket is approved. Until approval happens, all policies modified within a ticket are locked to that ticket and cannot be changed by anyone else.

Procedure

Step 1 Choose **System** (⚙) > **Change Management Workflow**, or click the **Ticket** (📋) shortcut menu.

Step 2 Click **Submit for Approval** (⬆) or **More** (⋮) > **Submit for Approval** for the open ticket.

Step 3 Optionally, enter a comment for your action.

Step 4 Click **Submit**.

The ticket is submitted for review and approval. Once approved, the changes made within the ticket are applied and the policies are unlocked.

Discard a ticket

Discard a ticket to remove any changes that were made within the ticket when you no longer need to proceed with those changes.

You cannot undo this action and retrieve a ticket with its changes. If you need them, you must create a new ticket and start over.

You cannot discard a ticket after you submit it. However, if the approver rejects the ticket, you can then discard it.

Note If you have permission to modify tickets, you can discard tickets that belong to another user. This makes it possible to deal with situations where an admin is on vacation or otherwise unavailable to manage an in-process ticket. If you have Review Ticket permission, you can alternatively reassign or take over the ticket rather than discard it.

Procedure

Step 1 Choose **System** (⚙) > **Change Management Workflow**, or click the **Ticket** (📋) shortcut menu.

Step 2 Click **Discard** (🗑) or **More** (⋮) > **Discard** for the ticket.

Step 3 Optionally, enter a comment for your action.

Step 4 Click **Discard**.

The ticket is discarded and any changes made within the ticket are removed.

Approving or rejecting a ticket

Approve or reject tickets to control which configuration changes become active and available for deployment in your network.

When a user submits a ticket, it must be approved for the changes made within the ticket to become active and available for deployment.

Whether you can approve your own tickets, or there is a separate approver, depends on your workplace policies and how user roles are assigned, not on the management software.

The Details view includes a summary of how many approvers are required for the ticket, and who has approved the ticket.

If the changes are inadequate or undesirable, you can reject the ticket. Rejected tickets go back to the submitter, who can then make additional changes and resubmit the ticket, or simply discard the ticket and the configuration changes it contains.

Procedure

Step 1 Choose **System** (⚙) > **Change Management Workflow**, or click the **Ticket** (📋) shortcut menu.

Step 2 On the **Review** tab, click **Preview** (🔍) for the ticket and evaluate the proposed changes.

You can also click **Validate** (🔍) or **More** (⋮) > **Validate** to check for errors.

Step 3 After completing your evaluation, do one of the following:

- To approve the ticket, click **Approve** (✓) or **More** (⋮) > **Approve**.
- To disapprove the ticket, click **Reject** (✗) or **More** (⋮) > **Reject**.

Step 4 Optionally, enter a comment for your action.

Step 5 Click **Approve** or **Reject**, as appropriate.

The ticket is approved or rejected based on your selection. Approved tickets become active and available for deployment. Rejected tickets return to the submitter for revision or disposal.

Take over or reassign tickets

Take over or reassign tickets to ensure continuous workflow management when the original ticket owner is unavailable or when ticket ownership needs to be transferred for deployment purposes.

Sometimes it might be necessary to take over a ticket that someone else created. For example, the ticket owner might be on vacation or is otherwise unavailable, and the ticket is blocking updates that need to be deployed.

You can also use this procedure to reassign your own ticket to another person.

Before you begin

Following are the permissions required to take over tickets:

- Admin user—You can assign tickets to yourself or other users.
- Modify or Review ticket + System > User Management > Users (custom role)—You can assign tickets to yourself or other users.

However, you can assign a user only if the user has the same role as the current ticket owner, or the Admin role. This ensures that the new user has the permissions required to configure the features currently modified within the ticket.

You cannot reassign a ticket that has been submitted for approval.

Follow these steps to take over or reassign tickets:

Procedure

Step 1 Choose **System** (⚙) > **Change Management Workflow**.

Step 2 Click on **All** under **Tickets in System**.

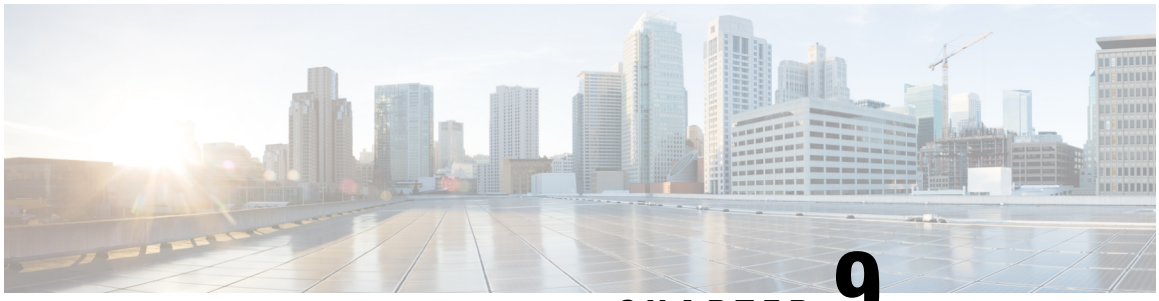
Step 3 Click **More** (⋮) > **Take Over Ticket**.

Step 4 Select the user who should now own the ticket.

The list of users is restricted to those who have the permissions required to edit the policies already changed within the ticket. For example, if a ticket contains changes to the access control policy, the list of users contains only those users who are allowed to modify the access control policy.

Step 5 Enter an optional comment and click **Takeover**.

The ticket ownership is transferred to the selected user, and the new owner can now manage and modify the ticket according to their permissions.



CHAPTER 9

CLI Users

Managed devices include a default **admin** account for CLI access. This chapter discusses how to create custom user accounts.

- CLI users, on page 181
- Guidelines for CLI users, on page 183
- Add an internal user at the CLI, on page 183
- Troubleshoot LDAP authentication connections, on page 186

CLI users

A CLI user is a user account that

- can be added as internal users or as external users on a LDAP or RADIUS server,
- is maintained separately on each managed device, and
- requires separate configuration for access to the Cloud-Delivered Firewall Management Center versus managed devices.

User account separation

When you add a user to the Cloud-Delivered Firewall Management Center, that user only has access to the Cloud-Delivered Firewall Management Center; you cannot then use that username to log directly into a managed device. You must separately add a user on the managed device.

Internal and external users

Managed devices support two categories of users—internal and external—for user access control.

Internal users authenticate through a local database on the device. External users authenticate through LDAP or RADIUS servers when they are not found in the local database. Both user types enable secure access management for network devices.

User authentication types

Managed devices support these user authentication methods:

- Internal user: The device checks a local database for user authentication.

- External user: If the user is not present in the local database, the system queries an external LDAP or RADIUS authentication server.

CLI access

CLI access is a command-line interface capability that

- provides a Firepower CLI running on top of Linux for device administration
- enables creation of internal users on devices using the CLI, and
- supports establishment of external users on Firewall Threat Defense devices using the Cloud-Delivered Firewall Management Center.

Security considerations for CLI access

Users with CLI Config level access can access the Linux shell using the **expert** command, and obtain `sudoers` privileges in the Linux shell, which can present a security risk.



Caution Users with CLI Config level access can access the Linux shell using the **expert** command, and obtain `sudoers` privileges in the Linux shell, which can present a security risk. For system security reasons, we strongly recommend:

- Only use the Linux shell under TAC supervision or when explicitly instructed by Firepower user documentation.
- Make sure that you restrict the list of users with CLI access appropriately.
- When granting CLI access privileges, restrict the list of users with Config level access.
- Do not add users directly in the Linux shell; only use the procedures in this chapter.
- Do not access Firepower devices using CLI expert mode unless directed by Cisco TAC or by explicit instructions in the Firepower user documentation.

CLI user roles

A CLI user role is an access control mechanism that

- determines the level of command access a user has on managed devices,
- controls whether users can execute configuration or non-configuration commands, and
- provides security by restricting unauthorized access to critical system functions.

Role types and permissions

On managed devices, user access to commands in the CLI depends on the role you assign.

- **None**: The user cannot log into the device on the command line.

- **Config**: The user can access all commands, including configuration commands. Exercise caution in assigning this level of access to users.
- **Basic**: The user can access non-configuration commands only. Allowed commands are dig, ping, and traceroute. Only internal users and Firewall Threat Defense external RADIUS users support the Basic role.

Guidelines for CLI users

This reference provides essential guidelines for CLI users, including username requirements, default account information, and user account limits for different device models.

Usernames

- You cannot add the same username for both internal and external users. If the external server uses a duplicate username, the deployment to the device fails.
- The username must be Linux-valid:
  - Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
  - All lowercase
  - Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Defaults

All devices include an **admin** user as a local user account; you cannot delete the **admin** user. The default initial password is **Admin123**; the system forces you to change this during the initialization process. The agent stores the modified credentials in the Windows registry and uses the same encryption as the Cloud-Delivered Firewall Management Center, which is Cipher Block Chaining (CBC). See the getting started guide for your model for more information about system initialization.

Number of user accounts

You can create a maximum of 43 user accounts for the Firepower 1000.

Add an internal user at the CLI

Create internal users to provide authorized access to the Firewall Threat Defense system with appropriate privilege levels.

Use the CLI to create internal users on the Firewall Threat Defense.

Procedure

Step 1 Log into the device CLI using an account with Config privileges.

The **admin** user account has the required privileges, but any account with Config privileges will work. You can use an SSH session or the Console port.

For certain Firewall Threat Defense models, the Console port puts you into the FXOS CLI. Use the **connect ftd** command to get to the Firewall Threat Defense CLI.

Step 2 Create the user account.

configure user add *username* {**basic** | **config**}

- *username*—Sets the username. The username must be Linux-valid:
  - Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
  - All lowercase
  - Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)
- **basic**—Gives the user basic access. This role does not allow the user to enter configuration commands. Allowed commands are dig, ping, and traceroute.
- **config**—Gives the user configuration access. This role gives the user full administrator rights to all commands.

Example:

This example adds a user account named johncrichton with Config access rights. The password is not shown as you type it.

```
> configure user add johncrichton config
Enter new password for user johncrichton: newpassword
Confirm new password for user johncrichton: newpassword
> show user
Login              UID   Auth Access  Enabled Reset    Exp Warn  Str Lock Max
admin             1000  Local Config  Enabled    No  Never  N/A  Dis   No N/A
johncrichton      1001  Local Config  Enabled    No  Never  N/A  Dis   No   5
```

Note
Tell users they can change their own passwords using the **configure password** command.

Step 3 (Optional) Adjust the characteristics of the account to meet your security requirements.

You can use these commands to change the default account behavior.

- **configure user aging** *username max_days warn_days*

  Sets an expiration date for the user's password. Specify the maximum number of days for the password to be valid followed by the number of days before expiration the user will be warned about the upcoming expiration. Both values are 1 to 9999, but the warning days must be less than the maximum days. When you create the account, there is no expiration date for the password.

- **configure user forcereset** *username*

  Forces the user to change the password on the next login.

- **configure user maxfailedlogins** *username number*

Sets the maximum number of consecutive failed logins you will allow before locking the account, from 1 to 9999. Use the **configure user unlock** command to unlock accounts. The default for new accounts is 5 consecutive failed logins. The admin account is exempt from being locked out after a maximum number of failed logins unless you enable security certification compliance.

- **configure user minpasswdlen** *username number*

  Sets a minimum password length, which can be from 1 to 127.

- **configure user strengthcheck** *username* {**enable** | **disable**}

  Enables or disables password strength checking, which requires a user to meet specific password criteria when changing their password. When a user's password expires or if the **configure user forcereset** command is used, this requirement is automatically enabled the next time the user logs in.

Step 4 Manage user accounts as necessary.

Users can get locked out of their accounts, or you might need to remove accounts or fix other issues. Use these commands to manage the user accounts on the system.

- **configure user access** *username* {**basic** | **config**}

  Changes the privileges for a user account.

- **configure user delete** *username*

  Deletes the specified account.

- **configure user disable** *username*

  Disables the specified account without deleting it. The user cannot log in until you enable the account.

- **configure user enable** *username*

  Enables the specified account.

- **configure user password** *username*

  Changes the password for the specified user. Users should normally change their own password using the **configure password** command.

  Caution
  Do not use the Linux **passwd** command in Expert Mode to change the admin user password. This command can cause file system corruption. Only use the regular Firewall Threat Defense CLI **configure user password admin** command (if you are not admin) or **configure password** command (if you are admin). If you do not know the password and can't log in at all, refer to the password recovery procedure.

- **configure user unlock** *username*

  Unlocks a user account that was locked due to exceeding the maximum number of consecutive failed login attempts.

Note
Firewall Threat Defense uses the Pluggable Authentication Module (PAM) framework for user authentication. Within this framework, the pam_unix.so module handles password hashing and verification using the SHA-512 algorithm

Passwords are hashed using the SHA-512 algorithm.

Troubleshoot LDAP authentication connections

This reference provides troubleshooting information for resolving LDAP authentication connection problems in your system configuration.

If you create an LDAP authentication object and it either does not succeed in connecting to the server you select or does not retrieve the list of users you want, you can tune the settings in the object.

If the connection fails when you test it, try the following suggestions to troubleshoot your configuration:

- Use the messages displayed at the top of the web interface screen and in the test output to determine which areas of the object are causing the issue.
- Check that the user name and password you used for the object are valid:
  - Check that you have the rights to browse to the directory indicated in your base-distinguished name by connecting to the LDAP server using a third-party LDAP browser.
  - Check that the user name is unique to the directory information tree for the LDAP server.
  - If you see an LDAP bind error 49 in the test output, the user binding for the user failed. Try authenticating to the server through a third-party application to see if the binding fails through that connection as well.
  - When configuring LDAP authentication, avoid using a backslash ("\") in the binding password. The presence of a backslash in the password causes the LDAP binding process to fail, resulting in unsuccessful external authentication.
- Check that you have correctly identified the server:
  - Check that the server IP address or host name is correct.
  - Check that you have TCP/IP access from your local appliance to the authentication server where you want to connect.
  - Check that access to the server is not blocked by a firewall and that the port you have configured in the object is open.
  - If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used for the server.
  - Check that you have not used an IPv6 address for the server connection if you are authenticating CLI access.
  - If you used server type defaults, check that you have the correct server type and click **Set Defaults** again to reset the default values.
- If you typed in your base-distinguished name, click **Fetch DNs** to retrieve all the available base distinguished names on the server, and select the name from the list.
- If you are using any filters, access attributes, or advanced settings, check that each is valid and typed correctly.

- If you are using any filters, access attributes, or advanced settings, try removing each setting and testing the object without it.
- If you are using a base filter or a CLI access filter, make sure that the filter is enclosed in parentheses and that you are using a valid comparison operator (maximum 450 characters, including the enclosing parentheses).
- To test a more restricted base filter, try setting it to the base distinguished name for the user to retrieve just that user.
- If you are using an encrypted connection:
  - Check that the name of the LDAP server in the certificate matches the host name that you use to connect.
  - Check that you have not used an IPv6 address with an encrypted server connection.
- If you are using a test user, make sure that the user name and password are typed correctly.
- If you are using a test user, remove the user credentials and test the object.
- Test the query that you are using by connecting to the LDAP server and using this syntax:

```
ldapsearch -x -b 'base_distinguished_name'
-h LDAPserver_ip_address -p port -v -D
'user_distinguished_name' -W 'base_filter'
```

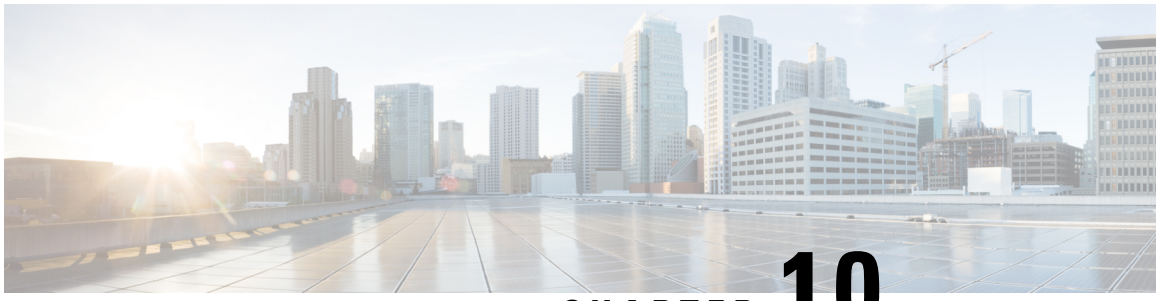
  For example, if you are trying to connect to the security domain on `myrtle.example.com` using the `domainadmin@myrtle.example.com` user and a base filter of (`cn=*`), you could test the connection using this statement:

```
ldapsearch -x -b 'CN=security,DC=myrtle,DC=example,DC=com'
-h myrtle.example.com -p 389 -v -D
'domainadmin@myrtle.example.com' -W '(cn=*)'
```

If you can test your connection successfully but authentication does not work after you deploy a platform settings policy, check that authentication and the object you want to use are both enabled in the platform settings policy that is applied to the device.

If you connect successfully but want to adjust the list of users retrieved by your connection, you can add or change a base filter or CLI access filter or use a more restrictive or less restrictive base DN.

While authenticating a connection to Active Directory (AD) server, rarely the connection event log indicates blocked LDAP traffic although the connection to AD server is successful. This incorrect connection log occurs when the AD server sends a duplicate reset packet. The Firewall Threat Defense device identifies the second reset packet as part of a new connection request and logs the connection with Block action.



CHAPTER 10

Configuration Deployment

This chapter describes how to download configuration changes to one or more managed devices.

- Configuration deployment, on page 189
- Requirements and prerequisites for deployment management, on page 200
- Best practices for deploying configuration changes, on page 200
- Configuration deployment, on page 201
- Deployment management, on page 209

Configuration deployment

Configuration deployment is a management process that enables centralized control and distribution of device settings from Cloud-Delivered Firewall Management Center to managed devices.

Configuration changes that require deployment

Cloud-Delivered Firewall Management Center marks out-of-date policies with red status text that indicates how many of its targeted devices need a policy update. Redeploy the policy to the devices to clear this status.

Deployment required

You must deploy changes after these configuration updates:

- Modifying an access control policy: any changes to access control rules, the default action, policy targets, Security Intelligence filtering, advanced options including preprocessing, and so on.
- Modifying any of the policies that the access control policy invokes: the SSL policy, network analysis policies, intrusion policies, file policies, identity policies, or DNS policies.
- Changing any reusable object or configuration used in an access control policy or policies it invokes:
  - network, port, VLAN tag, URL, and geolocation objects
  - Security Intelligence lists and feeds
  - application filters or detectors
  - intrusion policy variable sets
  - file lists

    - decryption-related objects and security zones

- Updating intrusion rules (SRU/LSP) or the vulnerability database (VDB).
- Upgrading managed devices.

You can change some of these configurations from multiple places in the web interface. For example, you can modify security zones using the object manager (**Objects** > **Object Management**), but modifying an interface type in a device configuration (**Devices** > **Device Management**) can also change a zone and require a deployment.

Deployment not required

These updates do not require a deployment:

- Security Intelligence updates to feeds and the global Block or Do Not Block lists.
- URL filtering data updates.
- Geolocation database (GeoDB) updates.

Deployment preview

A deployment preview is a snapshot that provides all the policy and object changes to be deployed on the device.

The policy changes include the new policies, changes in the existing policies, and the deleted policies. The object changes include the added and modified objects which are used in policies. The unused object changes are not displayed because they are not deployed on the device.

- **Comparison View**: Displays a side-by-side comparison of all the device configuration changes made since your last deployment.
- **Advanced View**: Displays the pending CLI commands to be applied to the device.

For more information about viewing the deployment preview, refer to Deploy configuration changes, on page 202.

The preview shows all the default values, even when they are not altered, along with the other configured settings when an interface or a platform settings policy is added for the first time. Similarly, the high availability-related policies and default values for settings are shown, even when they are not altered, in the first preview after a high availability pair is configured or disrupted.

To view changes due to an auto rollback, refer to Edit deployment settings, on page 1192.

Unsupported features:

- Object additions and attribute changes are displayed in the preview only if the objects are associated with any device or interface. Object deletions are not displayed.
- Preview is not supported for the policies:
    - High availability
    - Network discovery
    - Network analysis

- Device settings

- User information at the rule level is not available for intrusion policies.
- The preview does not show the reordering of rules across policies.

  For DNS policies, reordered rules appear in the preview list as rule additions and deletions. For example, moving a rule from position 1 to position 3 in the rule order is displayed as if the rule was deleted from position 1 and added as a new rule in position 3. Similarly, when a rule is deleted, the rules under it are listed as edited rules as they have changed their positions. The changes are displayed in the final order in which they appear in the policy.
- Preview is not supported in these HA scenarios:
  - If a device was in standalone mode and if a chain is made, then an auto-deployment is triggered. For that particular job, preview is not supported. On hover over the **Preview** (), a message is displayed that it is a HA bootstrap deployment, and no preview is supported.
  - **Configuration groups** - Consider a flow in which a device was initially standalone. Subsequently, three deployments took place. In the fourth deployment, the device was a HA bootstrap deployment. After these, the user deploys devices 5, 6, and 7. The deployment 7 is an HA break deployment, and the user deploys devices 8, 9, and 10.

    In this flow, the preview between 3 and 5 is not supported because 4 was a HA deployment. Similarly, the preview between 8 and 3 is also not supported. Preview is supported only from 3 to 1, 7,6, 5, 4, and 10, 9, and 8.
  - If a device is broken (HA is broken) then the new device is considered as a fresh device.

Guidelines for selective policy deployment

The Cloud-Delivered Firewall Management Center allows you to select a specific policy within the list of all the changes on the device that are due for deployment and deploy only the selected policy.

This table describes the policies that suppot selective deployment.

Table 6: Policies that support selective deployment

Type	Policy
Access control	• Access control policies • Prefilter policies • QoS policies
Threat prevention and inspection	• Intrusion policies • Malware and file policies • DNS policies

Type	Policy
Identity and discovery	• SSL / decryption policies • Identity policies • Network discovery
Device or network connectivity	• NAT policies • Routing policies • VPN policies

Use the this table to understand the constraints and conditions for applying selective policy deployment.

Table 7: Limitations for selective deployment

Type	Description	Scenarios
Full deployment	Full deployment is necessary for specific deploy scenarios, and the Firewall Management Center does not support selective deployment in such scenarios. If you encounter an error in such scenarios, you may choose to proceed by selecting all the changes for deployment on the device.	Full deployment is required for: • The first deployment after you upgrade devices. • The first deployment after you restore a device. • The first deployment after modifications to a device's interface settings. • The first deployment after modifications in the virtual router settings.
Associated policy deployment	The Cloud-Delivered Firewall Management Center identifies interdependent policies which are interlinked. When one of the interlinked policies is selected, the remaining interlinked policies are automatically selected.	Scenarios wherein an associated policy is automatically selected: • When a new object is associated with an existing policy. • When an existing policy's object is modified. Scenarios wherein multiple policies are automatically selected: • When a new object is associated with an existing policy, and the same object is already associated with other policies, all the associated policies are automatically selected. • When a shared object is modified, all the associated policies are automatically selected.

Type	Description	Scenarios
Interdependent policy changes (shown using color-coded tags)	The Cloud-Delivered Firewall Management Center dynamically detects dependencies in-between policies, and between the shared objects and the policies. The interdependency of the objects or policies is shown using color-coded tags.	Scenarios wherein color-coded interdependent policies or objects are automatically selected: • When all the out-of-date policies have interdependent changes. For example, when an access control policy, an intrusion policy, and a NAT policy are out-of-date. Since access control policy and NAT policy share an object, all policies are selected together for deployment. • When all out-of-date policies share an object, and the object is modified.
Access Policy Group specifications	Access Policy Group policies are listed together in the preview window under **Access Policy Group** when you click **View** ( ).	The scenarios and the expected behavior for Access Policy Group policies are: • If the access control policy is out-of-date, all other out-of-date policies under this group, except file policy and intrusion policy, are selected when the access control policy is selected for deployment. However, if the access control policy is out-of-date, intrusion and file policies can be individually selected or deselected irrespective of whether the access control policy is selected or not, unless there are any dependent changes. For example, if a new intrusion policy is assigned to an access control rule, it indicates that there are dependent changes, then both the access control policy and the intrusion policy will be automatically selected when either of them is selected. • If no access control policy is out-of-date, other out-of-date policies in this group can be selected and deployed individually.

System username

The Cloud-Delivered Firewall Management Center displays the username as **system** for specific management operations. This allows administrators to distinguish system-initiated actions from those performed by users.

The **system** username is applied to these operations:

- Rollback
- Upgrade
- Firewall Threat Defense backup and restore

- SRU update
- LSP update
- VDB update

Auto-enabling of application detectors

When you perform application control but disable required detectors, the Cloud-Delivered Firewall Management Center automatically enables the appropriate system-provided detectors upon policy deployment. If no system-provided detector exists for the application, the Cloud-Delivered Firewall Management Center enables the most recently modified user-defined detector.

Asset rediscovery with network discovery policy changes

When you deploy changes to a network discovery policy, the system deletes and then rediscovers MAC address, TTL, and hops information from the network map for the hosts in your monitored networks. Also, the affected managed devices discard any discovery data that has not yet been sent to the Cloud-Delivered Firewall Management Center.

Snort restart scenarios

A Snort restart scenario is a situation that

- causes the traffic inspection engine, referred to as the `Snort process`, on a managed device to restart,
- interrupts inspection until the process resumesWhether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196., and
- may result in resource demands that cause a small number of packets to drop without inspection when you deploy.

Common snort restart triggers

Any of these scenarios cause the Snort process to restart:

- You deploy a specific configuration that requires the Snort process to restart. For more information, refer to Configurations that restart the snort process when deployed or activated, on page 198.
- You make a change that immediately restarts the Snort process. For more information, refer to Actions that immediately restart the snort process, on page 199.
- Traffic activates the currently deployed Automatic Application Bypass (AAB) configuration. For more information, refer Configure automatic application bypass, on page 1189.
- Enabling or disabling **Logging connection events to RAM disk** feature. For more information, refer the section **Log to Ramdisk** in Troubleshoot Drain of FMC Unprocessed Events.

Inspection interruption warnings during deployment

When you deploy configurations, the **Inspect Interruption** column in the deploy page indicates if a deployed configuration will restart the Snort process on the Firewall Threat Defense device. When the traffic inspection engine referred to as the `Snort process` restarts, inspection is interrupted until the process resumes. Whether traffic is interrupted or passes without inspection during the interruption depends on how the traffic is handled by the device. You can proceed with the deployment, cancel the deployment and modify the configuration, or delay the deployment until a time when deploying would have the least impact on your network.

During deployment, when the **Inspect Interruption** column indicates **Yes** and you expand the device configuration listing, an icon will highlight any configuration type that may restart the Snort process. Moving your pointer over the icon displays a message saying that deployment may interrupt traffic.

This table summarizes how the deploy page displays inspection interruption warnings.

Table 8: Inspection interruption indicators

Type	Inspect interruption	Description
Firewall Threat Defense	**Inspect Interruption** ( )Yes	At least one configuration would interrupt inspection on the device if deployed, and might interrupt traffic depending on how the device handles traffic. You can expand the device configuration listing for more information.
	--	Deployed configurations will not interrupt traffic on the device.
	Undetermined	The system cannot determine if a deployed configuration may interrupt traffic on the device. Undetermined status is displayed before the first deployment after a software upgrade, or in some cases during a Support call.
	Errors ( )	The system cannot determine the status due to an internal error. Cancel the operation and click **Deploy** again to allow the system to redetermine the **Inspect Interruption** status. If the problem persists, contact Support.
sensor	--	The device identified as *sensor* is not the Firewall Threat Defense device; the system does not determine if a deployed configuration may interrupt traffic on this device.

For information on all configurations that restart the Snort process for all device types, see Configurations that restart the snort process when deployed or activated, on page 198.

Inspect traffic during policy apply

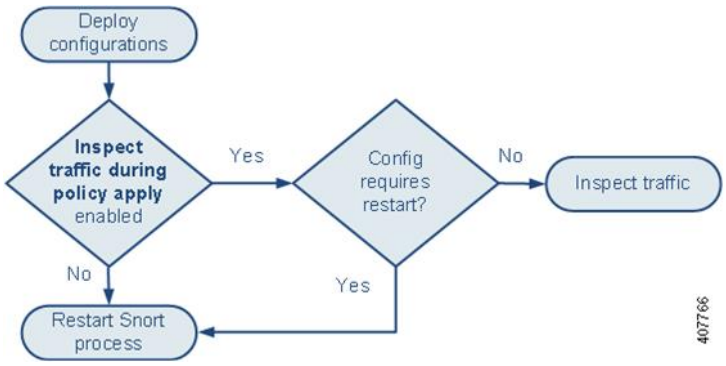
Inspect traffic during policy apply is an advanced access control policy general setting that enables managed devices to inspect traffic while configuration changes are applied, and allows inspection unless a deployed configuration requires the Snort process to restart.

Configuration options

You can configure this setting as follows:

- Enabled — Traffic is inspected during the deployment unless certain configurations require the Snort process to restart. When a Snort restart is not needed, the system first uses the currently deployed access control policy to inspect traffic, then switches to the policy being deployed.
- Disabled — Traffic is not inspected during the deployment. The Snort process always restarts when you deploy.

This graphic illustrates how Snort restarts can occur when you enable or disable **Inspect traffic during policy apply**.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. Refer to Snort restart traffic behavior, on page 196 and Configurations that restart the snort process when deployed or activated, on page 198.

Snort restart traffic behavior

This reference explains how different devices handle traffic when the Snort process restarts.

Table 9: The Firewall Threat Defense and the Firewall Threat Defense Virtual Restart Traffic Effects

Interface configuration	Restart traffic behavior
inline: **Snort Fail Open: Down**: disabled	dropped

Interface configuration	Restart traffic behavior
inline: **Snort Fail Open: Down**: enabled	passed without inspection Some packets can be delayed in buffer for several seconds before the system recognizes that Snort is down. This delay can vary depending upon the load distribution. However, the buffered packets are eventually passed.
routed, transparent (including EtherChannel, redundant, subinterface): **preserve-connection** enabled (**configure snort preserve-connection enable**; default) For more information, see *Cisco Secure Firewall Threat Defense Command Reference*.	existing TCP/UDP flows: passed without inspection so long as at least one packet arrives while Snort is down new TCP/UDP flows and all non-TCP/UDP flows: dropped Note that traffic drops even when **preserve-connection** is enabled for: • plaintext, passthrough prefilter tunnel traffic that matches an **Analyze** rule action or an **Analyze all tunnel traffic** default policy action • connections that do not match an access control rule and are instead handled by the default action • decrypted TLS/SSL traffic • a safe search flow • a captive portal flow
routed, transparent (including EtherChannel, redundant, subinterface): **preserve-connection** disabled (**configure snort preserve-connection disable**)	dropped
inline: tap mode	egress packet immediately, copy bypasses Snort
passive	uninterrupted, not inspected

Note In addition to traffic handling when the Snort process is down while it restarts, traffic can also pass without inspection or drop when the Snort process is busy, depending on the configuration of the Snort Fail Open **Busy** option (see Configure Inline Sets, on page 1128). A device supports either the Failsafe option or the Snort Fail Open option, but not both.

Note When the Snort process is busy but not down during configuration deployment, some packets may drop on routed, switched, or transparent interfaces if the total CPU load exceeds 60 percent.

Warning Do not reboot the system while the Snort Rule Update is in progress.

Snort-busy drops happen when snort is not able to process the packets fast enough. Lina does not know whether Snort is busy due to processing delay, or if is stuck or due to call blocking. When transmission queue is full, snort-busy drops occur. Based on Transmission queue utilization, Lina will try to access if the queue is being serviced smoothly.

Configurations that restart the snort process when deployed or activated

Deploying any of these configurations except AAB restarts the Snort process as described. Deploying AAB does not cause a restart, but excessive packet latency activates the currently deployed AAB configuration, causing a partial restart of the Snort process.

Access control policy advanced settings

- Deploy when **Inspect Traffic During Policy Apply** is disabled.
- Add or remove an SSL policy.

File policy

Deploy the first or last of any one of these configurations; note that while otherwise deploying these file policy configurations does not cause a restart, deploying non-file-policy configurations can cause restarts.

- Take either of these actions:
  - Enable or disable **Inspect Archives** when the deployed access control policy includes at least one file policy.
  - Add the first or remove the last file policy rule when **Inspect Archives** is enabled (note that at least one rule is required for **Inspect Archives** to be meaningful).
- Enable or disable **Store files** in a **Detect Files** or **Block Files** rule.
- Add the first or remove the last active file rule that combines the **Malware Cloud Lookup** or **Block Malware** rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**).

Note that access control rules that deploy these file policy configurations to security zones or tunnel zones cause a restart only when your configuration meets these conditions:

- Source or destination security zones in your access control rule must match the security zones associated with interfaces on the target devices.
- Unless the destination zone in you access control rule is *any*, a source tunnel zone in the rule must match a tunnel zone assigned to a tunnel rule in the prefilter policy.

Identity policy

- When SSL decryption is disabled (that is, when the access control policy does not include an SSL policy), add the first or remove the last active authentication rule.

An active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

Network discovery

- Enable or disable non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy.

Device management

- MTU: Change the highest MTU value among all non-management interfaces on a device.
- Automatic Application Bypass (AAB): The currently deployed AAB configuration activates when a malfunction of the Snort process or a device misconfiguration causes a single packet to use an excessive amount of processing time. The result is a partial restart of the Snort process to alleviate extremely high latency or prevent a complete traffic stall. This partial restart causes a few packets to pass without inspection, or drop, depending on how the device handles traffic.

TLS server identity discovery

Enabling or disabling **TLS Server Identity Discovery** will result in a Snort 2 restart.

Note Snort 3 does not restart when **TLS Server Identity Discovery** is enabled or disabled.

Updates

- System update: Deploy configurations the first time after a software update that includes a new version of the Snort binary or data acquisition library (DAQ).
- For managed devices running Snort 3, deploying configurations the first time after installing a vulnerability database (VDB) update may temporarily interrupt application detection, but there will be no traffic interruptions.

Actions that immediately restart the snort process

These actions immediately restart the Snort process without requiring deployment. How the restart affects traffic depends on how the target device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Applications or application detectors

- Activate or deactivate a system or custom application detector.
- Delete an activated custom detector.
- **Save and Reactivate** an activated custom detector.
- Create a user-defined application.

When you perform any of these actions, a message warns you that continuing will restart the Snort process. The message allows you to cancel. The system restarts Snort on all managed devices in the current domain and any child domain.

High availability configuration

Create or break a Firewall Threat Defense high availability pair.

When you create or break a high availability pair, a message warns you that creating or breaking a high availability pair restarts the Snort process on the primary and secondary devices. The message allows you to cancel.

Requirements and prerequisites for deployment management

Before using the configuration deployment feature in the Cloud-Delivered Firewall Management Center, ensure you meet these requirements.

- You must have **Administrator**, **Network Admin**, or **Security Approver** privileges to access the deployment features.
- Before deploying policy changes, you can compare the differences between the current device configuration and the pending policy, ensuring accuracy before you finalize the deployment. To use this feature, you must have the appropriate access rights, valid licenses for the specific policy, and be in the correct domain context.
- You must enable REST API access to preview the changes before a deployment. To enable the REST API access, follow the steps in *Enabling REST API Access* in the *Cisco Secure Firewall Management Center Administration Guide*.

Best practices for deploying configuration changes

Reliable management connection

The management connection between the Cloud-Delivered Firewall Management Center and the device is a secure, TLS-1.3-encrypted communication channel between itself and the device.

To ensure reliable management connection, use a direct management path instead of routing management traffic through an additional encrypted tunnel, such as a Site-to-Site VPN. If the tunnel fails, you could lose management connectivity.

To maintain a reliable management connection, exclude management traffic from the VPN tunnel if it exits a VPN-terminating interface.

Caution Avoid routing a device's management connection through a VPN tunnel that terminates on the device itself. If a configuration update causes the tunnel to drop, you will lose the management connection and cannot recover the device configuration unless you connect directly to the device.

Maximum concurrent deployments

Avoid deploying to more than 25 percent of the maximum devices allowed for a Cloud-Delivered Firewall Management Center in the same deployment task. For example, for the FMCv300, the maximum deployment task size is 75 devices (25 percent of 300). Deploying concurrently to more devices can cause performance issues.

Deployment of shared policies

For best performance, deploy to devices that use the same policies. Create a separate deployment task for each group of devices that share policies.

Time to deploy and memory limitations

Deployment time depends on multiple factors, including:

- The configurations you send to the device. For example, if you dramatically increase the number of Security Intelligence entries you block, deployment can take longer.
- Device model and memory. On devices with less memory, deployment can take longer.

Keep your configurations within the capability of your devices. If you exceed the maximum number of rules or policies supported by a target device, Cloud-Delivered Firewall Management Center displays a warning. The maximum depends on a number of factors—not only memory and the number of processors on the device, but also on policy and rule complexity. For information on optimizing policies and rules, refer to Best practices for access control rules, on page 2058.

Use a maintenance window to lessen the impact of traffic interruptions

We *strongly* recommend you deploy in a maintenance window or at a time when interruptions will have the least impact.

- When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. Refer to Snort restart traffic behavior, on page 196 and Configurations that restart the snort process when deployed or activated, on page 198.
- For the Firewall Threat Defense devices, the **Inspect Interruption** column in the **Deploy** dialog warns you when deploying might interrupt traffic flow or inspection. You can proceed with, cancel, or delay deployment. For more information, refer to Inspection interruption warnings during deployment, on page 195.

Configuration deployment

Configuration deployment is the process of deploying configuration and policy updates to your Fireall Threat Defense devices after you make changes. This process allows you to apply changes to multiple components including device configurations, security policies, and network discovery settings.

Deployment components

Deploying updates applies to these components:

- Device and interface configurations
- Device-related policies, such as NAT, VPN, QoS, and platform settings policy

- Access control and related policies: DNS, file, identity, intrusion, network analysis, prefilter, SSL
- Network discovery policy
- Intrusion rule updates
- Configurations and objects associated with any of these elements

You can configure the system for automatic deployment by scheduling a deploy task or by setting the system to deploy when importing intrusion rule updates. Automating policy deployment is especially useful if you allow intrusion rule updates to modify system-provided base policies for intrusion and network analysis. Intrusion rule updates can also modify default values for the advanced preprocessing and performance options in your access control policies.

You can view deployment status in the **Notifications** pane.

Deploy configuration changes

Apply new or updated device configurations to managed devices.

After you change a device configuration, you must deploy the changes for them to take effect. Deploy during a maintenance window or off-peak hours to minimize potential traffic and inspection disruption.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. Refer to Snort restart traffic behavior, on page 196 and Configurations that restart the snort process when deployed or activated, on page 198.

Before you begin

- Confirm that all managed devices use the same revision of the Security Zones object. If you have edited security zone objects: Do not deploy configuration changes to any device until you edit the zone setting for interfaces on *all* devices you want to sync. You must deploy to all managed devices at the same time.
- To preview the deployment changes, enable REST API access. To enable the REST API access, follow the steps in *Enabling REST API Access* in the *Cisco Secure Firewall Management Center Administration Guide* .

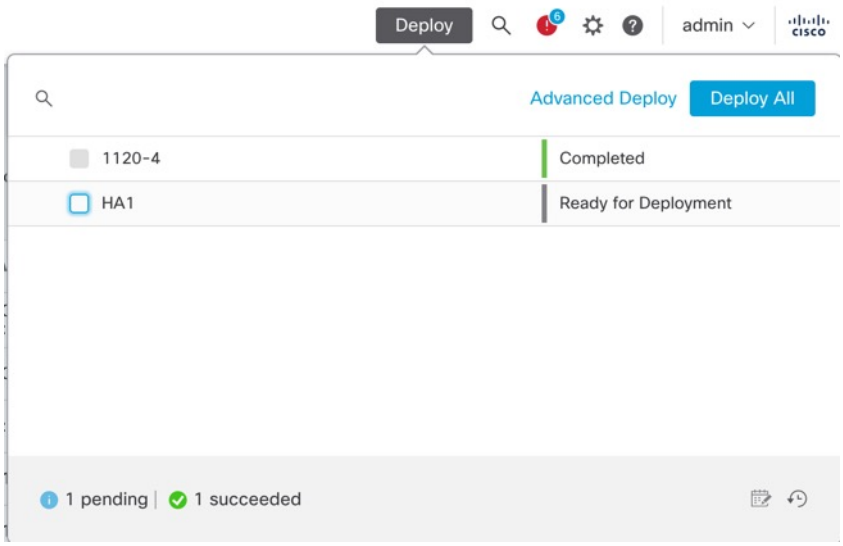
Note The deployment process fails if the device configuration is being read at the device CLI during deployment. Do not execute commands such as **show running-config** during the deployment.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center menu bar, click **Deploy**.

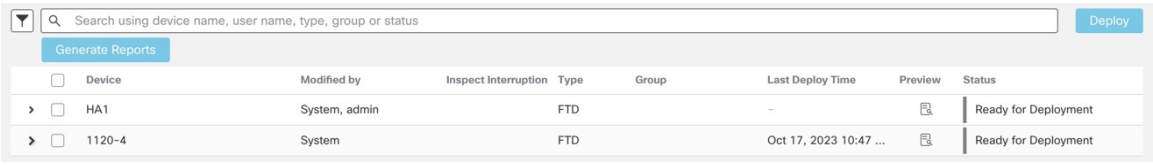
Step 2 For a quick deployment, check specific devices and then click **Deploy**, or click **Deploy All** to deploy to all devices. For additional deployment options, click **Advanced Deploy**.

Figure 13: Quick Deploy



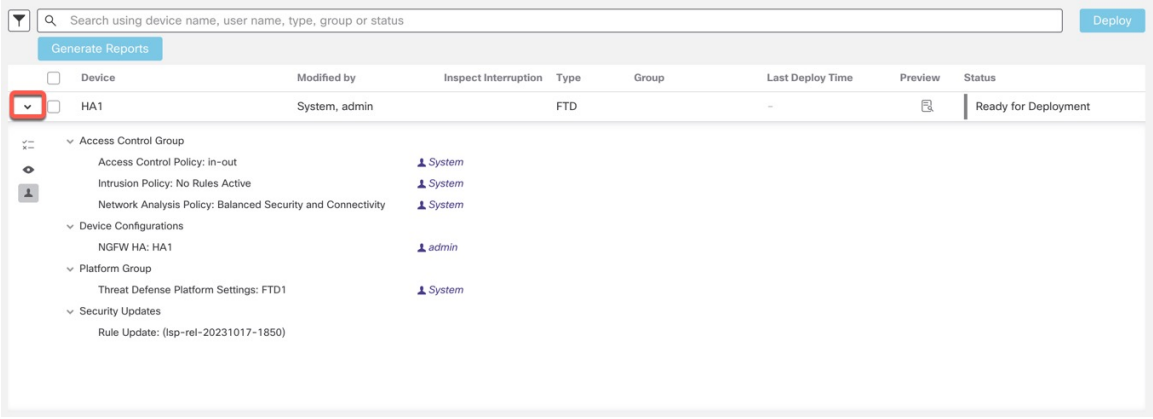
The rest of this procedure applies to the **Advanced Deploy** screen.

Figure 14: Advanced Deploy



Step 3 Click **Expand Arrow** (›) to view device-specific configuration changes to be deployed.

Figure 15: Expand



- The **Modified By** column lists the users who have modified the policies or objects. On expanding the device listing, you can view the users who have modified the policies against each policy listing. For

information about when the **System** user is shown (instead of the logged-in user), refer to System username, on page 193.

Note
Usernames are not provided for deleted policies and objects.

- The **Inspect Interruption** column indicates if traffic inspection interruption may be caused in the device during deployment.

  When the status indicates (Yes) that deploying will interrupt inspection, and perhaps traffic, on the Firewall Threat Defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** ( ).

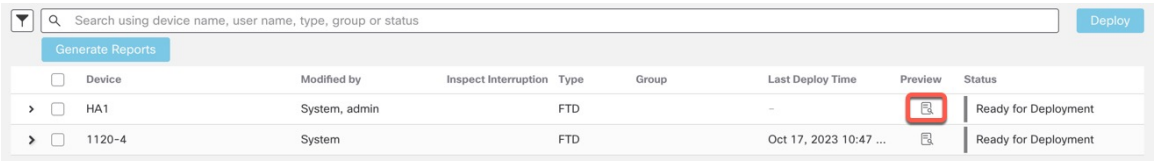
  If the entry is blank in this column for a device, then it indicates that there will be no traffic inspection interruptions on that device during deployment.

  For information about identifying configurations that interrupt traffic inspection and might interrupt traffic when deployed to the Firewall Threat Defense devices, refer to Inspection interruption warnings during deployment, on page 195.

- The **Last Modified Time** column specifies when you last made the configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.
- The **Status** column provides the status for each deployment. For more information, see Deployment status, on page 209.

Step 4 In the **Preview** column, click **Preview** ( ) to see the configuration changes that you can deploy.

Figure 16: Preview

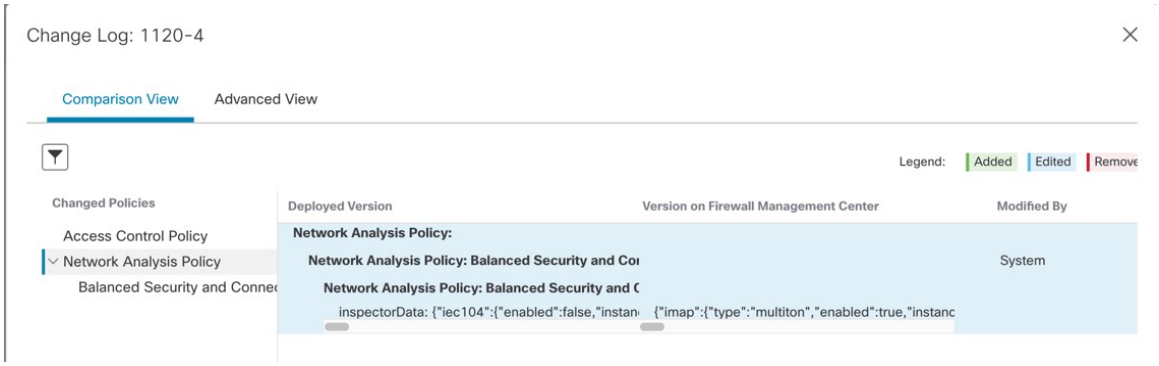


Note
If you change the Cloud-Delivered Firewall Management Center name in **System** ( ) > **Configuration** > **Information** , the deployment preview does not specify this change, yet it requires a deployment.

For unsupported features for Preview, see Deployment preview, on page 190.

The **Comparison View** tab lists all the policy and object changes. The left pane lists all the different policy types that have changed on the device, organized in a tree structure.

Figure 17: Comparison View



The **Filter**(lets you filter the policies at the user level and policy level.

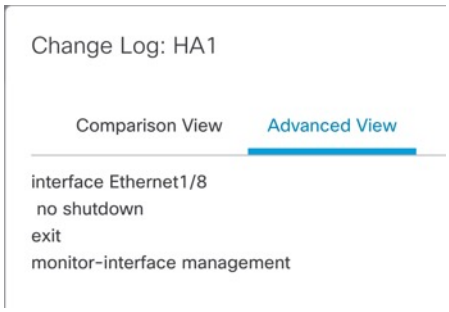
The right pane lists all the additions, changes, or deletions in the policy, or the object selected in the left pane. The two columns on the right pane provide the last deployed configuration settings (in the **Deployed Version** column) versus the changes that are due for deployment (in the **Version on Firewall Management Center** column). The last-deployed configuration settings are derived from a snapshot of the last saved deployment in the Cloud-Delivered Firewall Management Center and not from the device. The background colors of the settings are color-coded as per the legend available on the top-right of the page.

The **Modified By** column lists the users who have modified, or added the configuration settings. At the policy level, the Cloud-Delivered Firewall Management Center displays all the users who have modified the policy, and at the rule level, the Cloud-Delivered Firewall Management Center displays only the last user who has modified the rule.

You can download a copy of the change log by clicking the **Download Report** button.

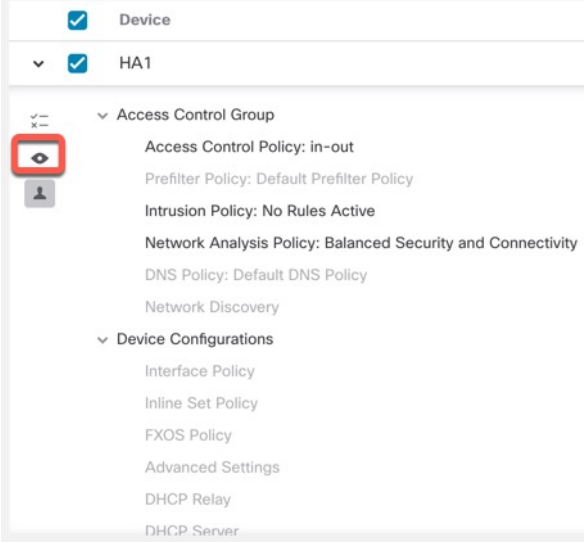
The **Advanced View** tab shows the CLI commands that will be applied. This view is useful if you are familiar with ASA CLI, which is used on the back end of the Firewall Threat Defense.

Figure 18: Advanced View



Step 5 Use **Show or Hide Policy** ( ) to selectively view or hide the associated unmodified policies.

Figure 19: Show or Hide Policy



Step 6 Check the box next to the device name to deploy all configuration changes, or click **Policy selection** ( ) to select individual policies or configurations to deploy while withholding the remaining changes without deploying them.

You can also view the interdependent changes for a certain policy or configuration using this option. The Cloud-Delivered Firewall Management Center dynamically detects dependencies between policies (for example, between an access control policy and an intrusion policy), and between the shared objects and the policies. Interdependent changes are indicated using color-coded tags to identify a set of interdependent deployment changes. When one of the deployment changes is selected, the interdependent changes are automatically selected.

For more information, refer to Guidelines for selective policy deployment, on page 191.

Note

- When the changes in shared objects are deployed, the impacted policies should also be deployed along with them. When you select a shared object during deployment, the impacted policies are automatically selected.
- Selective deployment is not supported for scheduled deployments and deployments using REST APIs. You can only opt for complete deployment of all the changes in these cases.
- The pre-deployment checks for warnings and errors are performed not only on the selected policies, but on all the policies that are out-of-date. Therefore, the warnings or errors list shows the deselected policies as well.
- Similarly, the **Inspect Interruption** column indication on the Deployment page considers all out-of-date policies and not just the selected policies. For information on the **Inspect Interruption** column, refer to Inspection interruption warnings during deployment, on page 195.

Step 7 After you select the devices or policies to deploy, click **Estimate** to get a rough estimate of the deployment duration .

Figure 20: Estimate

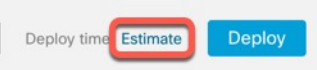
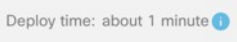


Figure 21: Deploy Time



The time duration is a rough estimate (having around 70% accuracy), and the actual time taken for deployment may vary for a few scenarios. The estimate is dependable for deployments of up to 20 devices.

When an estimate is not available, it indicates that the data is not available, since the first successful deployment on the selected device is pending. This situation could occur after the Cloud-Delivered Firewall Management Center reimage, version upgrade, or after a high availability failover.

Note
The estimate is incorrect and unreliable for bulk policy changes (in case of bulk policy migrations), and selective deployments because the estimate is based on the heuristic technique.

Step 8 For more details about the Pre-readiness assessment states, refer to .

Step 9 Click **Deploy**.

Configuration changes are deployed to the selected devices. You receive confirmation messages regarding success, warnings, or errors.

What to do next

If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- Deploy—Continue deploying without resolving warning conditions. Check the **Ignore warnings** checkbox, to ignore warnings and deploy the changes. You cannot proceed if the system identifies errors.
- Close—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.
- (Optional) Monitor deployment status; see *Viewing Deployment Messages* in the *Cisco Secure Firewall Management Center Administration Guide* .
- If the deployment fails, refer to Best practices for deploying configuration changes, on page 200 .
- During deployment, if there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For example, in a cluster environment, an erroneous configuration of an IP address that is not in the same subnet as the Site IPs is configured on the interface. Due to this error, deployment fails and the device attempts to clear the configuration while the rollback operation is being processed. These events collectively lead to a deployment failure that interrupts the traffic.

  See the following table to know what configuration changes may cause traffic interruption when deployment fails.

Table 10: Traffic interruption scenarios

Configuration Changes	Exists?	Traffic Impacted?
Threat Defense Service changes in an access control policy	Yes	Yes
VRF	Yes	Yes
Interface	Yes	Yes
QoS	Yes	Yes

Note The configuration changes interrupting traffic during deployment is valid only if both the Cloud-Delivered Firewall Management Center and the Firewall Threat Defense are of version 6.2.3 or later.

Redeploy existing configurations to a device

Force-deploying existing configurations to a managed device ensures that its policy rules are fully regenerated and applied. This procedure is recommended during maintenance windows or when disruption to network traffic is minimal. Force-deploy takes longer than regular deployment because it generates all policy rules for the device.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. Refer to Snort restart traffic behavior, on page 196 and Configurations that restart the snort process when deployed or activated, on page 198.

Before you begin

- Review the guidelines described in Best practices for deploying configuration changes, on page 200.
- Plan to perform the redeployment during a maintenance window or low-impact period.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** (✎) next to the device where you want to force deployment.

Step 3 Click **Device**.

Step 4 Click **Edit** (✎) next to the **General** section heading.

Step 5 Click **Force Deploy** ( ).

Step 6 Click **Deploy**.

The system identifies any errors or warnings with the configurations you are deploying. You can click **Proceed** to continue without resolving warning conditions. However, you cannot proceed if the system identifies an error.

What to do next

- (Optional) Monitor deployment status; see *Viewing Deployment Messages* in the *Cisco Secure Firewall Management Center Administration Guide*.
- If deploy fails, see Best practices for deploying configuration changes, on page 200.

Deployment management

Cloud-Delivered Firewall Management Center provides a set of deployement management options to oversee configuration changes across your Firewall Threat Defense devices. These options enable you to monitor real-time deployment status, track historical records, and ensure configuration integrity through rollback capabilities. Additionally, you can utilize built-in reporting and policy comparison features to audit changes, verify compliance, and troubleshoot issues.

Deployment status

On the **Deployment** page, the **Status** column provides the deployment status for each device. If a deployment is in progress, a live progress indicator is displayed; otherwise, one of these deployment status values appears:

- **Pending**: Indicates that there are changes in the device that need be deployed.
- **Warnings or errors**: Indicates that the pre-deployment checks have identified warnings or errors for the deployment. You may continue deployment if there are warnings, but not if there are errors.

Note The status column provides the warning or error status only for a single user session on the **Deployment** page. If you navigate away or refresh the page, the status changes to **Pending**.

- **Failed**: Indicates that the previous deployment attempt failed. Click the status to view the details.
- **In queue**: Indicates that deployment is initiated, and the system is yet to start the deployment process.
- **Completed**: Indicates that deployment has completed successfully.

View deployment history

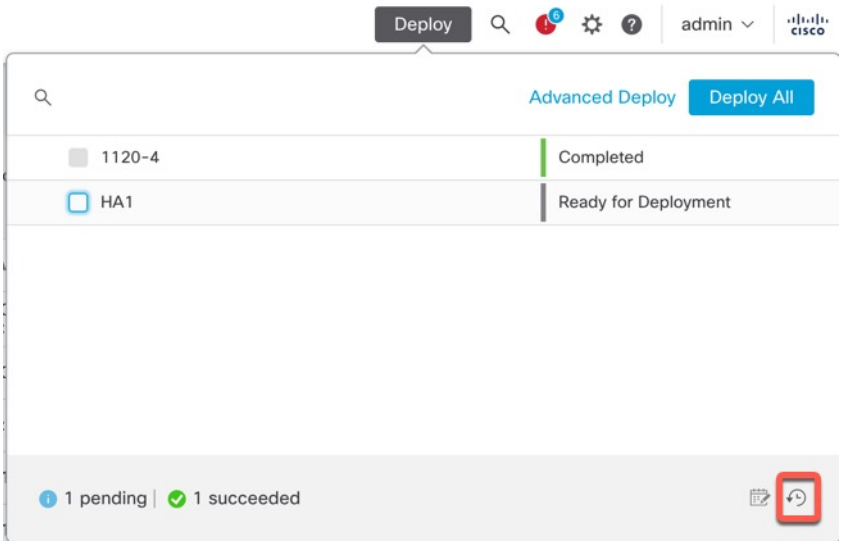
View deployment history to track and analyze past deployment activities, including successful deployments, failed deployments, and rollback operations.

In the deployment history, the last 10 successful deployments, the last 5 failed deployments, and last 5 rollback deployments are captured.

Procedure

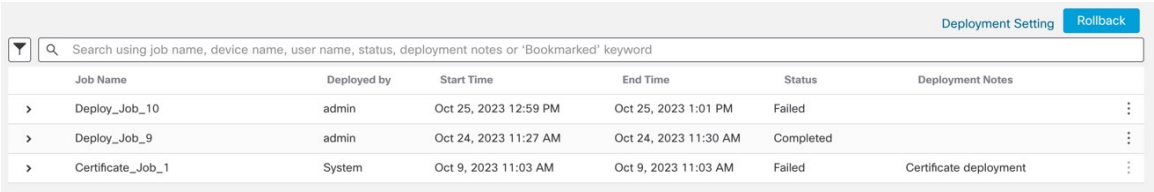
Step 1 On the Cloud-Delivered Firewall Management Center menu bar, click **Deploy** and then click **Deployment History** ( ).

Figure 22: Deployment history icon



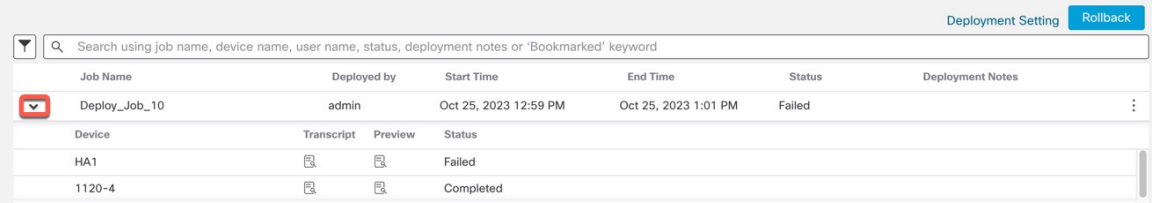
A list of all the previous deployment and rollback jobs is displayed in reverse chronological order.

Figure 23: Deployment history page



Step 2 Click **Expand Arrow** ( ) next to the required deployment job to view the devices included in the job and their deployment statuses.

Figure 24: Expand



- View notes in the **Deployment Notes** column.

  Deployment notes are custom notes that a user can add as part of the deployment, and these notes are optional.

Step 3 (Optional) Click **Transcript Details** ( ) to view the commands sent to the device, and the responses received.

Figure 25: Transcript details icon

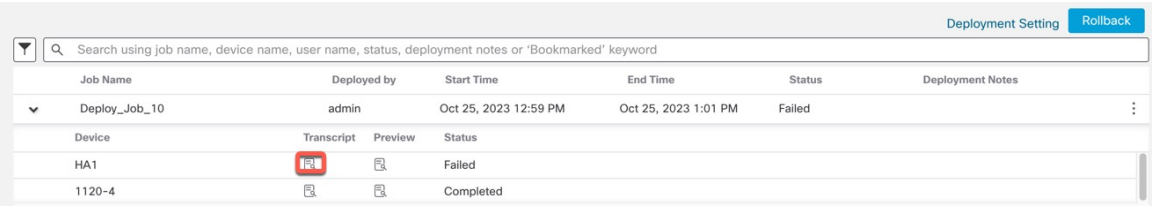
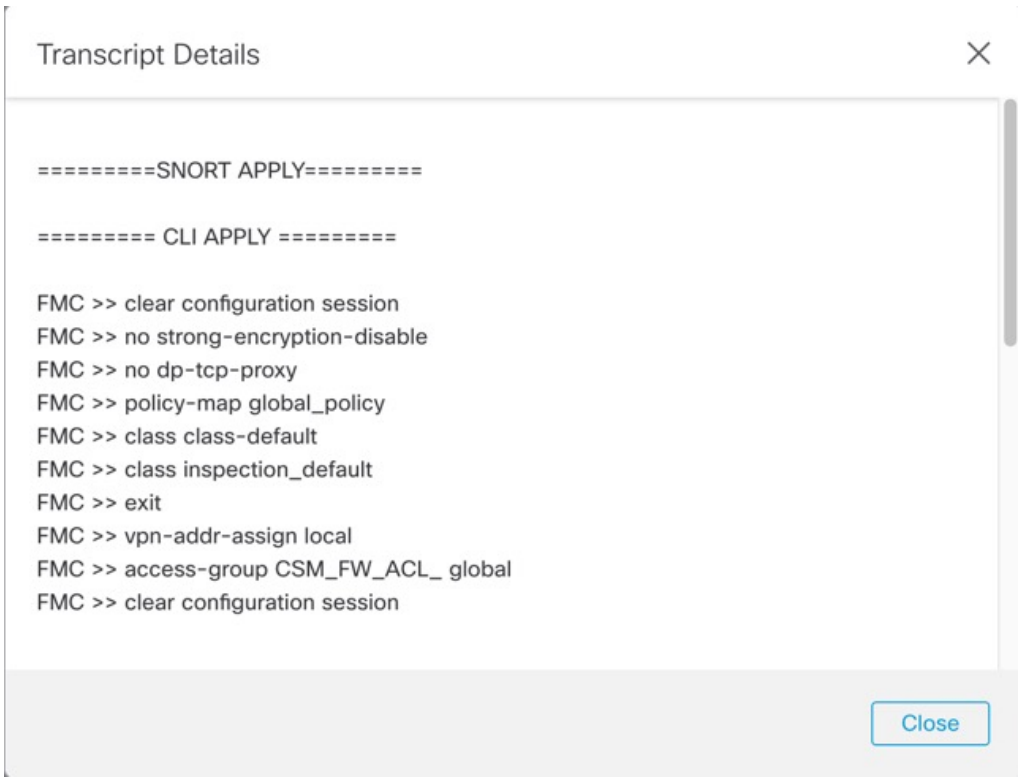


Figure 26: Transcript details



The transcript includes the following sections:

- **Snort Apply**—If there are any failures or responses from Snort-related policies, then the messages are displayed in this section. Normally, the section is empty.
- **CLI Apply**—This section covers features that are configured using commands that are sent to the device.
- **Infrastructure Messages**—This section shows the status of different deployment modules.

In the **CLI Apply** section, the deployment transcript includes commands that are sent to the device, and any responses returned from the device. These responses can be informative messages or error messages. For failed deployments, look for messages that indicate errors with the commands. Examining these errors can be particularly helpful if you are using FlexConfig policies to configure customized features. These errors can help you correct the script in the FlexConfig object that is trying to configure the commands.

Note
There is no distinction that is made in the transcript between commands that are sent for managed features and those generated from FlexConfig policies.

For example, the following sequence shows that Cloud-Delivered Firewall Management Center sent commands to configure GigabitEthernet0/0 with the logical name **outside**. The device responded that it automatically set the security level to 0. Firewall Threat Defense does not use the security level for anything.

```
========= CLI APPLY =========

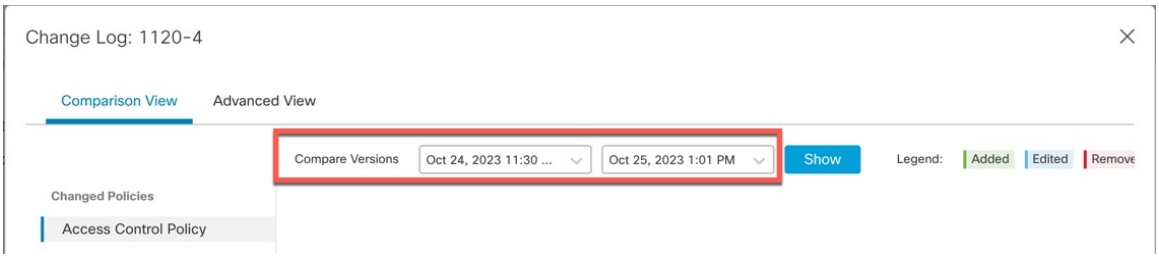
FMC >> interface GigabitEthernet0/0
FMC >>  nameif outside
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

Step 4 (Optional) Click **Preview** () to view the policy and object changes deployed on the device versus the previously deployed version.

Figure 27: Preview icon

a. To compare any two versions and view the change log, choose the required versions in the drop-down boxes and click the **Show** button. The drop-down boxes show the deployment job name and the end time of the deployment.

Figure 28: Compare versions



Note
The drop-down boxes also show failed deployments.

b. The **Modified By** column lists the users who have modified the policies or objects.

1. At the policy level, Cloud-Delivered Firewall Management Center displays all the user names who have modified the policy.
2. At the rule level, Cloud-Delivered Firewall Management Center displays the last user who has modified the rule.

c. You can also download a copy of the change log by clicking **Download Report**.

Note

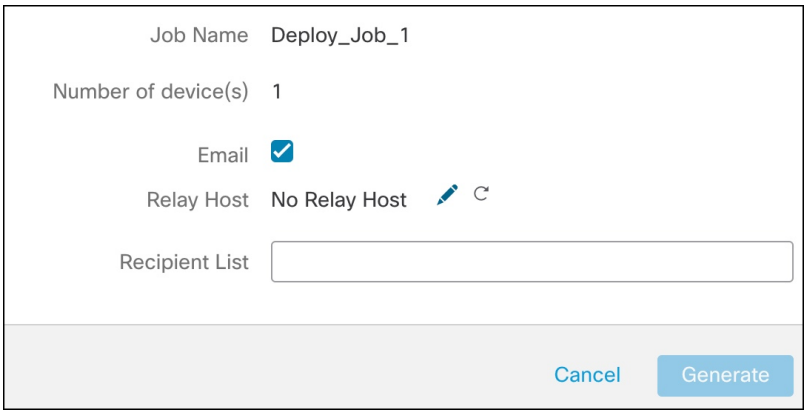
- Deployment history preview is not supported for certificate enrollments, HA operations, and failed deployments.
- When a device is registered, preview is not supported for the job history record that is created.

Step 5 (Optional) Against each deployment job, click the **More** (⋮) icon and execute other actions:

- **Bookmark**—To bookmark the deployment job.
- **Edit Deployment Notes**—To edit your custom deployment notes that you added for a deployment job.
- **Generate Report**—To generate a deployment report, which can be used for auditing. This report includes job properties with preview and transcript information, and the report can be downloaded as a PDF file.

a. Click **Generate Report** to generate a deployment report.

Figure 29: Generate report



b. In the **Generate Report** popup window, check the **Email** checkbox.

c. In the **Recipient List**, you can enter multiple email adddresses, separated by semicolons.

d. Click **Generate** to generate the report, and this report is emailed to the recipients.

e. In the Notifications task tab, you can track the progress. After the report generation is complete, click the link in the notification task tab to download the PDF report.

Deployment rollback

Deployment rollback is the process of reverting a device to a previously saved configuration state. This operation allows you to restore the device to its earlier state if a recent policy deployment results in unintended traffic behavior.



Note Rollback is a disruptive operation. Executing a rollback terminates all existing connections and resets routing tables, which will cause a temporary interruption of traffic.

Identifying the disruptive configuration

When a deployment has caused unintended traffic interruption, identify the specific change responsible for the condition. To compare current and previous deployments:

Before a rollback

1. Choose **Deploy** > **Deployment History**, expand the last deployed job that caused the traffic disruption, and click the **Preview** ().

   The preview page lets you compare deployments to identify specific changes from previous configurations.

2. After identifying the change causing the problem, rectify the configuration and redeploy it on the device.

After a rollback

1. After a successful rollback, choose **Deploy** > **Deployment**, and click the **Preview** icon next to the rolled back device.
2. View the changes between the rolled back configuration and the current changes in the Cloud-Delivered Firewall Management Center that are pending deployment.
3. After identifying the change causing the problem, rectify the configuration, and redeploy it on the device.

Rollback reverts most configurations, but some exceptions apply. Refer the table below for details.

Rollback guidelines

Follow these guidelines to ensure successful configuration rollbacks.

- You have to perform a deployment before you can roll back again.
- After you perform a rollback, the rolled back devices are marked as out-of-date on the Cloud-Delivered Firewall Management Center. The changes you made to the configuration are still pending for the next deployment. To see the pending changes, choose **Deploy** > **Deployment**, and click the **Preview** icon next to the rolled back device.

- For devices with very large access lists, if the **Object Group Search** setting is disabled, then the rollback operation may take a longer duration to complete. To verify the **Object Group Search** setting, choose **Devices** > **Device Management**, and then select the device and click **Edit Advanced Settings**.
- For the Firepower 4100/9300, make sure your current Firewall Chassis Manager interface configuration is the same for any rollback versions. Otherwise, the rollback interface configuration may not match your actual interfaces.
- Independent certificate enrollments are also listed as deployment jobs in the Deployment History page. However, you cannot roll back to these versions. A rollback from a deployment version created after certificate enrollments reverts the certificate associations as well. In the next deployment after a rollback, manually associate the certificates before proceeding with the deployment.
- If a deployment for a device with a FlexConfig object configured with a deployment frequency set to **Once** is rolled back, then you will not be able to redeploy the object even though it is displayed as out-of-date on the Preview page. After a rollback, you will have to manually unassign and then reassign the FlexConfig object to the device before the next deployment.

Rollback limitations

Observe these limitations that restrict the use of rollback functionality.

- You can roll back to any one of the last 10 versions before the currently deployed version. Rollback to versions prior to these are not supported. The rollback icon is greyed out for unsupported versions.
- Rollback is not supported if the manager access interface (Manager or data interface) is different between the rollback version and the current version.
- If you upgrade the Cloud-Delivered Firewall Management Center, all rollback versions from the previous software release will no longer be available for devices, even if you did not upgrade the devices.
- If you upgrade the device, you can only roll back to versions on the current software release.
- For High Availability, rollback is not supported in these scenarios:
  - When the version you want to roll back to contains the high-availability bootstrap configuration. In other words, the deployment when you first formed high availability for the standalone devices.
  - When a device that is currently in standalone mode was part of a high availability pair in the previous deployment version.
- For clustering:
  - Rollback is not supported when a device that is currently in standalone mode was part of a cluster in the previous deployment version.
  - (Secure Firewall 3100/4200 and Firewall Threat Defense Virtual in a private cloud) If you change the clustering bootstrap configuration or add or delete nodes, you cannot roll back to a version prior to those changes.

Table 11: Configurations reverted and not reverted during rollback

Configurations that are reverted during a rollback	Configurations that are not reverted during a rollback
• All policy configurations • Interface configurations • SRU configurations • VDB configurations • LSP configurations • VPN configurations • FXOS configurations	• Snort binaries • Geo DB

Perform a rollback

Use rollback to revert a device to a previously deployed configuration when the traffic through the device is affected in an unintended way after a policy deployment.

You can roll back a device to a previously deployed configuration. After a policy deployment, if the traffic through the device is affected in an unintended way, rollback provides an option to revert the device to the earlier state, which existed before the faulty deployment.

Rollback reverts the configuration only on the selected devices.

Procedure

Step 1 Choose **Deploy >Deployment History** ( ).

A list of all the previous deployment jobs is displayed in reverse chronological order.

Step 2 Click **Rollback**.

Step 3 Filter the list of devices displayed by either clicking **Job** and choosing a job from the **Selected Job** drop-down list, or by clicking **Device List**.

Step 4 (Optional) Enter the device name in the **Search Device** search box to filter the device list.

Step 5 Check the box next to the devices you want to roll back and choose the version for each device from the **Rollback Version** drop-down list.

Figure 30: Selected job list

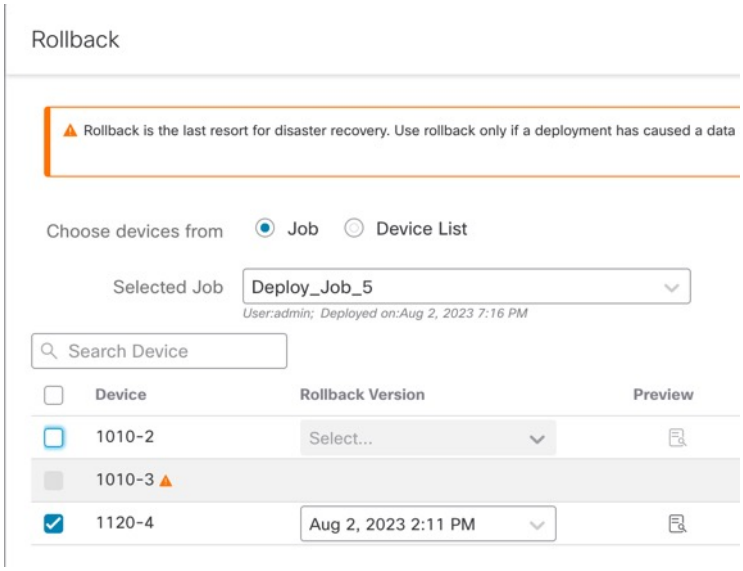
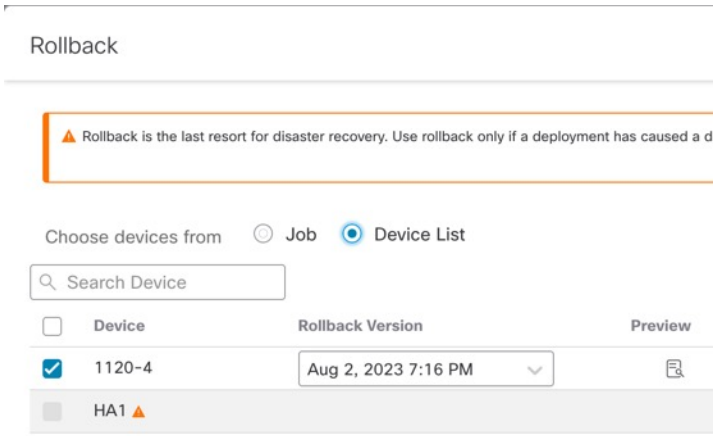


Figure 31: Device list



The job names and associated deployment notes are also listed for a particular rollback version.

Step 6 (Optional) Click the **Preview** ( ) to view the changes deployed in the selected version.

Step 7 Click **Rollback** to restore the selected devices to the chosen configuration version.

What to do next

To verify rollback status, go to **Deploy** > **Deployment** and check the rollback status next to the device name.

View the deployment rollback transcript

View the CLI commands that were executed during a successful rollback operation to understand what changes were made to the device configuration.

Rollback transcript is a written version of the commands that are sent to the device, along with the responses returned from the device. If a rollback operation has failed, the reason for the failure can be viewed by clicking **Deploy**, and then click the **Deployment History** (). However, to know the CLI commands executed for a successful rollback operation, follow the steps given below after a rollback operation has completed.

Note

- The CLI command information is available after a rollback is complete and is available only until the next deployment. The first deployment after a rollback operation erases all the rollback-related information.
- For any rollback deployment, the deployment notes are updated automatically as rollback job. In the **Deployment History** page, the user can filter the rollback jobs easily using the search option.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center menu bar, choose **System** () > **Health** > **Monitor**.
Step 2 Select the device, which was rolled back, from the left pane.
Step 3 Click **View System & Troubleshooting Details** link.
Step 4 Click **Advanced Troubleshooting**.
Step 5 Click **Threat Defense CLI**.
Step 6 Select **show** from the **Command** drop-down box.
Step 7 Enter `running` in the **Parameter** field.
Step 8 Click **Execute**.

Download policy changes report for multiple devices

Download reports on the policy and object changes made since your last deployment for multiple Firewall Threat Defense devices. Reports help you review pending changes before completing deployment. The reports are provided as a zip file containing:

- A pending changes report for each device, previewing modifications to policy or objects that are to be deployed on the device. For more information, refer to Deploy configuration changes, on page 202 and Deployment Preview.
- A consolidated report categorizing each device by report status.

Procedure

Step 1 Choose **Deploy** > **Advanced Deploy**.
Step 2 Check the check box next to the devices for which you want to generate a pending policy changes report, and then click **Pending Changes Reports**.
Step 3 Click **Pending Changes Reports**. Reports are generated in the background.

Step 4 On the Cloud-Delivered Firewall Management Center menu bar, choose **Notifications** > **Tasks** to view the report generation task.

When the report request task is complete, the download link appears within the task notification.

Step 5 Click the **Download Report(s)** link to download the reports.

Compare policies

Policy comparison allows you to examine the differences between two policies or between a saved policy and the running configuration to review policy changes for compliance with your organization's standards or to optimize system performance.

You can compare these policy types:

- DNS
- File
- Health
- Identity
- Network Analysis
- SSL

The comparison view displays both policies in a side-by-side format. Differences between the two policies are highlighted:

- Blue indicates that the highlighted setting is different in the two policies, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy but not the other.

Before you begin

You can compare policies only if you have access rights and any required licenses for the specific policy, and you are in the correct domain for configuring the policy.

Procedure

Step 1 Access the management page for the policy you want to compare:

- DNS—**Policies** > **Access Control heading** > **DNS**
- File—**Policies** > **Access Control heading** > **Malware & File**
- Health—**System** (⚙) > **Health** > **Policy**
- Identity—**Policies** > **Access Control heading** > **Identity**
- Network Analysis—**Policies** > **Access Control heading** > **Access Control**, and then click **Network Analysis Policy** or **Policies** > **Access Control heading** > **Intrusion**, and then click **Network Analysis Policies**.

Note

If your custom user role limits access to the first path listed here, use the second path to access the policy.

- SSL—**Policies** > **Access Control heading** > **Decryption**

Step 2 Click **Compare Policies**.

Step 3 From the **Compare Against** drop-down list, choose the type of comparison you want to make:

- To compare two different policies, choose **Other Policy**.
- To compare two revisions of the same policy, choose **Other Revision**.
- To compare another policy to the currently active policy, choose **Running Configuration**.

Step 4 Depending on the comparison type you choose, you have various choices:

- If you are comparing two different policies, choose the policies you want to compare from the **Policy A** and **Policy B** drop-down lists.
- If you are comparing the running configuration to another policy, choose the second policy from the **Policy B** drop-down list.

Step 5 Click **OK**.

Step 6 Review the comparison results:

- Comparison Viewer—To use the comparison viewer to navigate individually through policy differences, click **Previous** or **Next** above the title bar.
- Comparison Report—To generate a PDF report that lists the differences between the two policies, click **Comparison Report**.

Generate current policy reports

You can generate two types of reports for most policies: a single-policy report detailing the policy's current saved configuration, and a comparison report listing differences between two policies. You can generate a single-policy report for all policy types except health.

Note Intrusion policy reports combine the settings in the base policy with the settings of the policy layers, and make no distinction between which settings originated in the base policy or policy layer.

Before you begin

- Ensure you have the required access rights and licenses for the specific policy.
- Verify you are operating within the correct domain for the policy you wish to report.

Procedure

Step 1 Access the management page for the policy for which you want to generate a report:

- Access Control—**Policies** > **Access Control heading** > **Access Control**

- DNS—**Policies** > **Access Control heading** > **DNS**
- File—**Policies** > **Access Control heading** > **Malware & File**
- Health—**System** (⚙) > **Health** > **Policy**
- Identity—**Policies** > **Access Control heading** > **Identity**
- Intrusion—**Policies** > **Access Control heading** > **Intrusion**
- NAT—**Devices** > **NAT**
- Network Analysis—**Policies** > **Access Control heading** > **Access Control**, and then click **Network Analysis Policy** or **Policies** > **Access Control heading** > **Intrusion**, and then click **Network Analysis Policies**

  Note
  If your custom user role limits access to the first path listed here, use the second path to access the policy.

- SSL—**Policies** > **Access Control heading** > **Decryption**

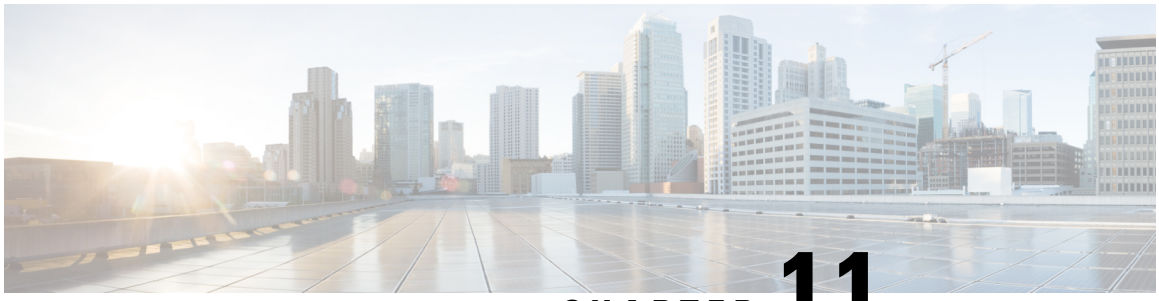
Step 2 Click **Report** (📄) next to the policy for which you want to generate a report.



PART III

System Settings

- System Configuration, on page 225
- Users, on page 243
- Updates, on page 263
- Licenses, on page 275
- Security certifications compliance, on page 315



CHAPTER 11

System Configuration

This chapter explains how to configure system configuration settings on the Cloud-Delivered Firewall Management Center.

- Meet requirements and prerequisites for the system configuration, on page 225
- Manage the Cloud-Delivered Firewall Management Center system configuration, on page 226
- Configure access control preferences, on page 226
- Audit logs, on page 227
- Audit log certificates, on page 231
- Change reconciliation, on page 237
- Email notifications, on page 239
- Intrusion policy preferences, on page 239
- Network analysis policy preferences, on page 241

Meet requirements and prerequisites for the system configuration

This reference lists the requirements and prerequisites for the system configuration, including supported models, domains, and user roles.

Model support

Firewall Management Center

Supported domains

Global

User roles

Admin

Manage the Cloud-Delivered Firewall Management Center system configuration

Configure the system to establish and maintain the essential settings for the Cloud-Delivered Firewall Management Center. Ensure that these settings align with organizational requirements.

The system configuration identifies basic settings for the Cloud-Delivered Firewall Management Center. Use this task when you need to review or update the fundamental configuration options for the management center.

Before you begin

Access to the Cloud-Delivered Firewall Management Center interface may be required. Follow these steps to manage the Cloud-Delivered Firewall Management Center system configuration.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Use the navigation panel to choose configurations to change.

When you complete the steps, the basic settings for the Cloud-Delivered Firewall Management Center will update to maintain proper system operation.

What to do next

Verify that the configuration changes are applied and the system is functioning as expected.

Configure access control preferences

You can track changes to access control rules by allowing or requiring users to comment when they save. This allows you to assess why critical policies in a deployment were modified. By default, this feature is disabled.

Configure object optimization to evaluate and optimize network or host policy objects that are used in rules. The system then creates associated network object groups on the device. For more information, refer to Object-group optimization.

Procedure

Step 1 Choose **System** (⚙) > **Configuration** > **Access Control Preferences**.

Step 2 From the **Comments on rule change** list, choose an option:

- **Disabled**: Disables comments for access control rule changes.
- **Optional**: Allows users to optionally add a comment when they change an access control rule.

- **Required**: Requires users to add a comment when they change an access control rule.

Step 3 From the **Object-group optimization** list, choose an option:

- **Enabled**: Enables object-group optimization. The setting takes effect after deployment.
- **Disabled**: Disables object-group optimization.

Warning
For the first deployment to a threat defense device after you enable object-group optimization, the device can take several minutes to an hour to reevaluate the policy configuration and optimize object groups. CPU utilization on the device can also increase. Schedule the deployment during a low-traffic period or a maintenance window.

Step 4 Click **Save**.

What to do next

Deploy the access control policies for object-group optimization to take effect.

Audit logs

An audit log is a security record that

- records user activity in read-only logs,
- allows review, sorting, filtering, deletion, and reporting of audit information, and
- supports streaming audit log messages to external servers such as syslog and HTTP servers.

Audit log review and streaming options

You can review audit log data in several ways:

- Use the web interface: .

  Audit logs appear in a standard event view. In this view, you can view, sort, and filter audit log messages based on any item. You can also delete audit information, generate reports, and view detailed reports of user changes.

- Stream audit log messages to the syslog: Stream audit logs to syslog, on page 228.
- Stream audit log messages to an HTTP server: Stream audit logs to an HTTP server, on page 230.

Streaming audit log data to an external server allows you to conserve space on the Cloud-Delivered Firewall Management Center. Note that sending audit information to an external URL may affect system performance.

Optionally, you can secure the channel for audit log streaming, enable TLS and mutual authentication using TLS certificates; refer to Audit log certificates, on page 231.

Streaming to multiple syslog servers

You can stream audit log data to a maximum of five syslog servers. However, if you have enabled TLS for secured audit log streaming, you can stream only to a single syslog server.

Streaming configuration changes to syslog

You can stream configuration changes as part of audit log data to syslog by specifying the configuration data format and the hosts. The Cloud-Delivered Firewall Management Center supports backup and restore of the audit configuration log.

Stream audit logs to syslog

Configure audit log streaming to syslog servers to ensure centralized monitoring and compliance tracking for your system. Enable external logging of configuration changes and security events.

When this feature is enabled, audit log records appear in the syslog in the following format:

```
Date Time Host: [Tag] Sender: User_Name@User_IP, Subsystem, Action
```

Where the local date, time, and originating hostname precede the bracketed optional tag, and the sending device name precedes the audit log message.

For example, if you specify a tag of `FMC-AUDIT-LOG` for audit log messages from your management center, a sample audit log message from your Cloud-Delivered Firewall Management Center could appear as follows:

```
Mar 01 14:45:24 localhost: [FMC-AUDIT-LOG] Dev-MC7000: admin@10.1.1.2, Operations >
Monitoring, Page View
```

If you specify a severity and facility, these values do not appear in syslog messages; instead, they tell the system that receives the syslog messages how to categorize them.

This task is relevant when you need to send audit logs from your management center to an external syslog server for compliance or operational monitoring. Use this procedure when integrating with centralized log management solutions.

Before you begin

Make sure the Cloud-Delivered Firewall Management Center can communicate with the syslog server. When you save your configuration, the system uses ICMP/ARP and TCP SYN packets to verify syslog server connectivity. By default, the system uses port 514/UDP to stream audit logs. To secure the channel, manually configure port 1470 for TCP.

Follow these steps to stream audit logs to a syslog server:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Audit Log**.

Step 3 Choose **Enabled** from the **Send Audit Log to Syslog** drop-down menu.

Step 4 The following fields are applicable only for audit logs sent to syslog:

Option	Description
Send Configuration Changes	To include configuration changes syslog in the audit log streaming, from the drop-down, select the relevant options: • JSON—the syslog includes detailed differences in the configuration changes. • API—the syslog includes API to retrieve the detailed differences in the configuration changes. • None—to have all other audit logs except details of the configuration changes.
Host	The IP address or the fully qualified name of the syslog server to which you will send audit logs. You can add a maximum of five syslog hosts, separated by commas. **Note** You can specify multiple syslog hosts, only when TLS is disabled for the Audit Server Certificate.
Facility	The subsystem that creates the message. Choose a facility described in SYSLOG alert facilities, on page 491. For example, choose AUDIT.
Severity	The severity of the message. Choose a severity described in Syslog severity levels, on page 492.
Tag	An optional tag to include in audit log syslog messages. Best practice: Enter a value in this field to easily differentiate audit log messages from other, similar syslog messages such as health alerts. For example, if you want all audit log records sent to the syslog to be labeled with `FMC-AUDIT-LOG`, enter `FMC-AUDIT-LOG` in the field.

Step 5 (Optional) To test whether the IP address of the syslog servers is valid, click **Test Syslog Server**.

The system sends the following packets to verify whether the syslog server is reachable:

a. ICMP echo request

b. TCP SYN on 443 and 80 ports

c. ICMP time stamp query

d. TCP SYN on random ports

Note
If the Cloud-Delivered Firewall Management Center and syslog server are in the same subnet, ARP is used instead of ICMP.

The system displays the result for each server.

Step 6 Click **Save**.

After you complete this task, audit logs are sent from the management center to the specified syslog server, allowing you to monitor and review system activity externally.

What to do next

Verify that the syslog server is receiving audit log messages as expected. Check for log entries corresponding to configuration changes or security events.

Stream audit logs to an HTTP server

Stream audit logs from the device to an HTTP server to centralize log management and support compliance requirements.

Use this task to forward audit logs from your device to an external HTTP server for monitoring, analysis, or compliance.

When this feature is enabled, the appliance sends audit log records to an HTTP server in the following format:

```
Date Time Host [Tag] Sender: User_Name@User_IP, Subsystem, Action
```

Where the local date, time, and originating hostname precede the bracketed optional tag, and the sending appliance name precedes the audit log message.

For example, if you specify a tag of `FROMMC`, a sample audit log message could appear as follows:

```
Mar 01 14:45:24 localhost [FROMMC] Dev-MC7000: admin@10.1.1.2, Operations > Monitoring,
Page View
```

Before you begin

Make sure the device can communicate with the HTTP server.

Follow these steps to stream audit logs to an HTTP server.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Audit Log**.

Step 3 Optionally, in the **Tag** field, enter the tag name that you want to appear with the message. For example, if you want all audit log records to be preceded with `FROMMC`, enter `FROMMC` in the field.

Step 4 Choose **Enabled** from the **Send Audit Log to HTTP Server** drop-down list.

Step 5 In the **URL to Post Audit** field, designate the URL where you want to send the audit information. Enter a URL that corresponds to a Listener program that expects the HTTP POST variables as listed:

- `subsystem`
- `actor`
- `event_type`
- `message`
- `action_source_ip`
- `action_destination_ip`

- `result`
- `time`
- `tag` (if defined; see Step 3)

Caution
To allow encrypted posts, use an HTTPS URL. Sending audit information to an external URL may affect system performance.

Step 6 Click **Save**.

When you complete this task, the device sends audit logs to the specified HTTP server. Centralized monitoring and record-keeping become possible.

Audit log certificates

An audit log certificate is a security credential that

- secures communications between the Cloud-Delivered Firewall Management Center and a trusted audit log server,
- enables authentication and encryption using Transport Layer Security (TLS), and
- supports mutual authentication and certificate revocation checks for enhanced security.

Client certificates

Client certificates are required to secure communications between the Cloud-Delivered Firewall Management Center and the audit log server.

- Generate a certificate signing request (CSR). Submit it to a Certificate Authority (CA) for signing. Import the signed certificate onto the Cloud-Delivered Firewall Management Center.
- Use the local system configuration: Obtain a signed audit log client certificate for the Cloud-Delivered Firewall Management Center, on page 233 and Import an audit log client certificate into the Cloud-Delivered Firewall Management Center, on page 234.

Server certificates

Server certificates are optional and provide additional security through mutual authentication.

- Require mutual authentication between the Cloud-Delivered Firewall Management Center and the audit log server by loading one or more certificate revocation lists (CRLs).
- You cannot stream audit logs to servers with revoked certificates listed in those CRLs.
- Secure Firewall supports CRLs encoded in Distinguished Encoding Rules (DER) format. These are the same CRLs used to validate HTTPS client certificates for the Cloud-Delivered Firewall Management Center web interface.
- Use the local system configuration: Require valid audit log server certificates, on page 235.

Audit log certificate example

For example, a signed client certificate imported onto the Cloud-Delivered Firewall Management Center enables secure TLS communication with an audit log server, ensuring log integrity and confidentiality.

Non-compliant audit log certificate

A certificate that is not signed by a trusted Certificate Authority or is listed in a certificate revocation list (CRL) cannot be used to stream audit logs securely.

Audit log certificate analogy

An audit log certificate is like a passport for secure communication, verifying identity and granting access between the management center and the audit log server.

Secure audit log streaming

Configure secure audit log streaming to ensure that audit logs are transmitted safely between the Cloud-Delivered Firewall Management Center and a trusted server.

When you stream the audit log to a trusted HTTP server or syslog server, use Transport Layer Security (TLS) certificates to secure the channel between the Cloud-Delivered Firewall Management Center and the server.

You must generate a unique client certificate for each appliance you want to audit. Mutual authentication requires the client certificate to be signed by the same CA as the server certificate. For more considerations on client and server certificate requirements, refer to Audit log certificates, on page 231.

Before you begin

Refer to ramifications of requiring client and server certificates at Audit log certificates, on page 231. Ensure you have access to a recognized certificate authority (CA) for signing client certificates.

Follow these steps to secure audit log streaming.

Procedure

Step 1 Obtain and install a signed client certificate on the Cloud-Delivered Firewall Management Center.

a) Obtain a signed audit log client certificate for the Cloud-Delivered Firewall Management Center, on page 233:

Generate a Certificate Signing Request (CSR) from the Cloud-Delivered Firewall Management Center based on your system information and the identification information you supply.

- Submit the CSR to a recognized, trusted certificate authority (CA) to request a signed client certificate.
- If you will require mutual authentication between the Cloud-Delivered Firewall Management Center and the audit log server, ensure the client certificate is signed by the same CA that signed the server certificate.

b) After you receive the signed certificate from the certificate authority, import it into the Cloud-Delivered Firewall Management Center. Refer to Import an audit log client certificate into the Cloud-Delivered Firewall Management Center, on page 234.

Step 2 Configure the communication channel with the server to use Transport Layer Security (TLS) and enable mutual authentication.

Step 3 If you have not yet configured audit log streaming, do so now.

Refer to Stream audit logs to syslog, on page 228 or Stream audit logs to an HTTP server, on page 230.

Audit logs are securely streamed to the trusted server using TLS, with mutual authentication enabled. The Cloud-Delivered Firewall Management Center and the server verify each other's certificates, ensuring secure transmission and integrity of audit log data.

What to do next

Monitor the audit log streaming status and verify certificate validity periodically. Renew client and server certificates before expiration to maintain secure log streaming.

Obtain a signed audit log client certificate for the Cloud-Delivered Firewall Management Center

Obtain a signed audit log client certificate for the Cloud-Delivered Firewall Management Center. This certificate enables secure communication between the appliance and the audit log server. It ensures that the server can authenticate audit log data and that the data is transmitted securely.

The system generates certificate request keys in Base-64 encoded PEM format.

Perform this task on the active appliance to generate a certificate signing request for audit log client authentication. Do not perform this task on a standby appliance in a high availability setup.

Before you begin

Keep this in mind:

- To ensure security, use a globally recognized and trusted certificate authority (CA) to sign your certificate.
- If you will require mutual authentication between the appliance and the audit log server, the same certificate authority must sign both the client certificate and the server certificate.

Follow these steps to obtain a signed audit log client certificate for the Cloud-Delivered Firewall Management Center:

Procedure

Step 1 Choose **System** (⚙) > **Configuration** and click **Audit Log Certificate**.

Step 2 Click **Generate New CSR**.

Step 3 Enter a country code in the **Country Name (two-letter code)** field.

Step 4 Enter a state or province postal abbreviation in the **State or Province** field.

Step 5 Enter a **Locality or City**, an **Organization** name, and an **Organizational Unit (Department)** name.

Step 6 Enter the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field.

Note
If the common name and the DNS hostname do not match, audit log streaming will fail.

Step 7 Click **Generate** and then open a new blank file with a text editor.

Step 8 Copy the block of text from the certificate request, including the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines. Paste this information into a blank text file.

Step 9 Save the file as *`clientname`*`.CSR`. Replace *`clientname`* with the appliance name where you plan to use the certificate. Click **Close**.

After completing these steps, you will have a certificate signing request (CSR) file that can be submitted to a trusted certificate authority. After it is signed, you can import the certificate to the appliance for secure audit log streaming.

What to do next

- Submit the certificate signing request to the certificate authority you selected, according to the guidelines in the 'Before You Begin' section.
- When you receive the signed certificate, import it to the appliance; refer to Import an audit log client certificate into the Cloud-Delivered Firewall Management Center, on page 234.

Import an audit log client certificate into the Cloud-Delivered Firewall Management Center

Importing an audit log client certificate enables secure communication between the Secure Firewall Management Center and external systems during audit log transmission. This task verifies that you import the correct signed certificate and any necessary intermediate certificates for secure operations.

In a Secure Firewall Management Center high availability setup, you must use the active peer when importing the audit log client certificate. This procedure is relevant when configuring audit log security or after obtaining a new signed client certificate.

Before you begin

- Obtain a signed audit log client certificate for the Cloud-Delivered Firewall Management Center, on page 233.
- Make sure you are importing the signed certificate for the correct Cloud-Delivered Firewall Management Center.
- If the signing authority requires an intermediate CA, ensure you provide the certificate chain. The same CA must sign both the client certificate and the intermediate certificates.

Follow these steps to import an audit log client certificate into the Secure Firewall Management Center.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **System** (⚙) > **Configuration**.

Step 2 Click **Audit Log Certificate**.

Step 3 Click **Import Audit Client Certificate**.

Step 4 Open the client certificate in a text editor, copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines. Paste this text into the **Client Certificate** field.

Step 5 To upload a private key, open the private key file and copy the entire block of text, including the `BEGIN RSA PRIVATE KEY` and `END RSA PRIVATE KEY` lines. Paste this text into the **PRIVATE KEY** field.

Step 6 Open any required intermediate certificates, copy the entire block of text for each, and paste it into the **Certificate Chain** field.

Step 7 Click **Save**.

The audit log client certificate and any required intermediate certificates are successfully imported into the Secure Firewall Management Center.

Require valid audit log server certificates

The system supports validating audit log server certificates using imported CRLs in Distinguished Encoding Rules (DER) format.

Note If you choose to verify certificates using CRLs, the system uses the same CRLs to validate both audit log server certificates and certificates used to secure the HTTP connection between an appliance and a web browser.

Before you begin

- Understand the ramifications of requiring mutual authentication and of using certificate revocation lists (CRLs) to ensure that certificates are still valid. Refer to Audit log certificates, on page 231.
- Obtain and import the client certificate following the steps in Secure audit log streaming, on page 232 and the topics referenced in that procedure.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center , choose **System** (⚙) > **Configuration** .

Step 2 Click **Audit Log Certificate** .

Step 3 To use Transport Layer Security to securely stream the audit log to an external server, select **Enable TLS** .

When TLS is enabled, the syslog client (Cloud-Delivered Firewall Management Center) verifies the certificate received from the server. The connection between the client and the server succeeds only if server certificate verification is successful. For this verification process, these conditions must be met:

- Configure the syslog server to send the certificate to the client.
- Add (import) a CA certificate to the client to verify the server certificate:
  - You must import the CA certificate during the import of the client certificate.
  - If the issuing CA is a subordinate CA, add the issuing CA before adding the signing CA from the subordinate CA (Root CA).

Step 4 If you do not want the client to authenticate itself against the server, but wish to accept the server certificate when the certificate is issued by the same CA (not recommended), complete these steps.

a) Deselect **Enable Mutual Authentication** .

Important
Ensure that the server is configured to trust the client without verifying any client certificates.

b) Click **Save** and skip the remainder of this procedure.

Step 5 (Optional) To enable client certificate verification by the audit log server, select **Enable Mutual Authentication**.

Important
The **Enable Mutual Authentication** option is applicable only when TLS is enabled.

When mutual authentication is enabled, the syslog client (Cloud-Delivered Firewall Management Center) sends a client certificate to the syslog server for verification. The client uses the same CA certificate of the CA who signed the server certificate of the syslog server. The connection succeeds only if client certificate verification is successful. For this verification process, these conditions must be met:

- Configure the syslog server to verify the certificate received from the client.
- Add a client certificate to be sent to the syslog server. This certificate must be signed by the same CA who signed the server certificate of the syslog server.

Note
To use mutual authentication for streaming Audit Log to the Syslog server, use PKCS#8 format for the private key instead of PKCS#1 format. Use this command line to convert PKCS#1 keys to PKCS#8 format:

```
openssl pkcs8 -topk8 -inform PEM -outform PEM
-nocrypt -in

 PKCS1 key file name

-out

 PKCS8 key filename
```

Step 6 (Optional) To automatically recognize server certificates that are no longer valid:

a) Select **Enable Fetching of CRL** .

Important
This option is displayed only when you select the **Enable Mutual Authentication** check box. However, the **Enable Fetching of CRL** option is applicable only when the TLS option is enabled. The use of CRL

is for server certification verification, and it is not dependent on the use of Mutual Authentication which is for enabling client certificate verification.

When you enable fetching of the CRL, the client creates a scheduled task to regularly update (download) the CRL or CRLs. The CRLs are used for server certificate verification. Verification fails if a CRL from the CA specifies that the server certificate has been revoked.

b) Enter a valid URL to an existing CRL file and click **Add CRL**.

Repeat to add up to 25 CRLs.

c) Click **Refresh CRL** to load the current CRL or CRLs from the specified URL or URLs.

Step 7 Verify that you have a valid server certificate generated by the same certificate authority that created the client certificate.

Step 8 Click **Save**.

What to do next

(Optional) Set the frequency of CRL updates. Refer to Schedule certificate revocation list (CRL) updates.

View the audit log client certificate on the Cloud-Delivered Firewall Management Center

View the audit log client certificate to confirm its presence and details on the appliance you are logged in to. This helps ensure that audit log operations are secured with the correct certificate.

You can view the audit log client certificate only for the appliance that you are logged in to.

Before you begin

Follow these steps to view the audit log client certificate on the appliance:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Audit Log Certificate**.

The appliance you are logged in to displays the audit log client certificate details. For high availability pairs, only the active peer shows the certificate.

Change reconciliation

A change reconciliation report is a configuration monitoring tool that

- captures snapshots whenever a user saves changes to the system configuration,
- combines information from these snapshots to present a clear summary of recent system changes, and

- provides a detailed report of changes made over the past 24 hours through email.

User change summaries

The change reconciliation report displays both the previous value for each configuration and the value after changes.

Users can view summaries of each distinct change in chronological order, beginning with the most recent.

- Lists both previous and updated values for each configuration.
- Summarizes multiple changes to the same configuration.
- Orders changes chronologically, starting with the most recent.

Change reconciliation report example

An example change reconciliation report includes a User section that lists configuration changes, showing both the previous and updated values for each item.

Non-reconciled changes

Changes made without reconciliation are not summarized or tracked in a report, making it difficult to monitor user activity.

Change reconciliation as a transaction log

Change reconciliation is similar to a transaction log, where each modification is recorded and summarized for review.

Configure change reconciliation

Configure change reconciliation to enable the system to generate and send daily reports of configuration changes, supporting audit and compliance requirements. This helps administrators monitor and review all changes made to the system within a 24-hour period.

Use this task when you need to track and review configuration changes made to your system for security, compliance, or operational awareness.

Change reconciliation is especially useful in environments where multiple administrators make changes or where audit trails are required for regulatory purposes. Perform this configuration after setting up your email server to ensure reports are delivered successfully.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Change Reconciliation**.

Step 3 Check the **Enable** check box.

Step 4 Choose the time of day you want the system to send out the change reconciliation report from the **Time to Run** drop-down lists.

Step 5 Enter email addresses in the **Email to** field.

Tip
Once you have added email addresses, click **Resend Last Report** to send recipients another copy of the most recent change reconciliation report.

Step 6 If you want to include policy changes, check the **Include Policy Configuration** check box.

Step 7 If you want to include all changes over the past 24 hours, check the **Show Full Change History** check box.

Step 8 Click **Save**.

After you complete this task, the system will automatically generate and email daily reports of all configuration changes, allowing you to monitor and review system modifications for compliance and troubleshooting.

Change reconciliation options

The change reconciliation options control whether the system includes records of policy changes and the full change history in the change reconciliation report. These options determine the scope and detail of reported changes.

The **Include Policy Configuration** option controls whether the system includes records of policy changes in the change reconciliation report. It includes changes to access control, intrusion, system, health, and network discovery policies. If you do not select this option, the report will not show changes to any policies. This option is available only on Cloud-Delivered Firewall Management Centers.

The **Show Full Change History** option controls whether the system includes records of all changes over the past 24 hours in the change reconciliation report. If you do not select this option, the report includes only a consolidated view of changes for each category.

Note The change reconciliation report does not include changes to Firewall Threat Defense interfaces and routing settings.

Email notifications

An email notification is a system alert mechanism that enables automated emailing of event-based, status, change reconciliation, and data-pruning reports.

Email notification restrictions for static mail relay hosts

You cannot configure a mail host. TThe mail relay host is hardcoded to a static host. It is set to email-smtp.us-west-2.amazonaws.com with authorization. For notifications, the email sender is set to cdo-alert@cisco.com.

Intrusion policy preferences

An intrusion policy preference is a configuration option that

- enables monitoring of critical policies,
- tracks changes to those policies, and
- supports management of policy settings in a deployment.

Set intrusion policy preferences

Configure intrusion policy preferences to manage how policy changes are tracked, logged, and monitored for security intelligence. You can enable or disable comments on policy changes, audit logging, user overrides for deleted Snort 3 rules, and Talos Threat Hunting Telemetry.

This task is relevant when you need to control how intrusion policy changes are documented, audited, and monitored for advanced threat detection and compliance.

Use this configuration to ensure that policy modifications are properly tracked and that security events are sent to Cisco Talos for analysis when required. Perform this task during initial setup or when updating security and compliance requirements for your deployment.

Before you begin

Follow these steps to set intrusion policy preferences:

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Intrusion Policy Preferences**.

Step 3 You have these options:

- **Comments on policy change**: Check this check box to track policy-related changes using the comment functionality when users modify intrusion policies. With policy change comments enabled, administrators can quickly assess why critical policies in a deployment were modified.

  If you enable comments on policy changes, you can make the comment optional or mandatory. The Cloud-Delivered Firewall Management Center prompts the user for a comment when each new change to a policy is saved.

- **Write changes in Intrusion Policy to audit log**: Check this check box to record the changes to the intrusion policies to the audit logs. This option is enabled by default.

- **Retain user overrides for deleted Snort 3 rules**: Check this check box to get notifications for changes to any *overridden* system-defined rules during LSP updates. When you enable this feature, the system retains the rule overrides in the new replacement rules that are added as part of the LSP update. On the Cloud-Delivered Firewall Management Center menu bar, click **Notifications** (message center). Then click **Tasks** to view the notifications. This option is enabled by default.

- **Talos Threat Hunting Telemetry**: Check this check box to allow Cisco Talos to conduct threat hunting and to gather critical security intelligence. When enabled, a special set of threat-hunting rules is added to the global intrusion policy. Although the threat-hunting rules are processed like regular IPS rules, events generated by these rules do not appear in the Cloud-Delivered Firewall Management Center's event tables. These events are sent to Talos as telemetry for analysis. This option is enabled by default.

  Note

- If you send firewall events to the Cisco Security Cloud via a direct connection by registering your Cloud-Delivered Firewall Management Center to the cloud tenancy using your Security Cloud Control account, your Security Cloud Control account must have a Security Analytics and Logging license in order to forward threat-hunting rule events to Talos.

After completing this task, your intrusion policy preferences are set according to your selections. Policy changes are tracked, logged, and, if enabled, threat-hunting telemetry is sent to Cisco Talos for analysis.

Network analysis policy preferences

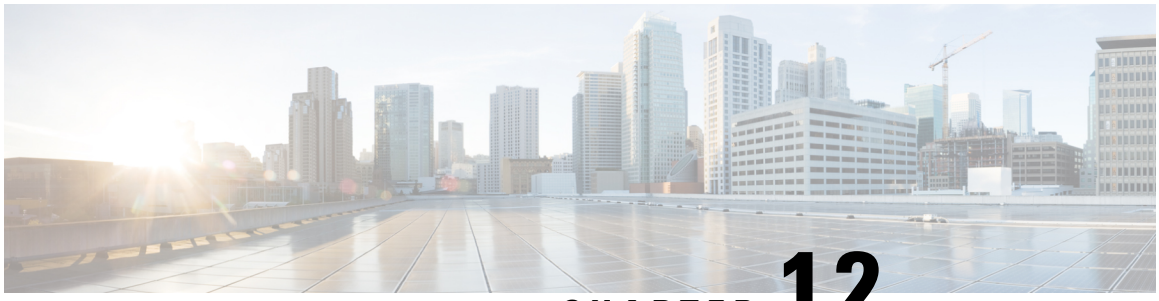
Network analysis policy preferences are configuration settings that enable tracking and documentation of policy modifications within the system.

Policy Change Tracking Options

You can configure the system to track policy-related changes using the comment functionality when users modify network analysis policies. With policy change comments enabled, administrators can quickly assess why critical policies in a deployment were modified.

If you enable comments on policy changes, you can make the comment optional or mandatory. The system prompts the user for a comment when each new change to a policy is saved.

Optionally, you can have changes to network analysis policies written to the audit log.



CHAPTER 12

Users

The Cloud-Delivered Firewall Management Center includes default **admin** accounts for web and CLI access. This chapter discusses how to create custom user accounts.

- CLI users, on page 243
- Create a User Record with Your Security Cloud Control Username, on page 246
- External authentication configuration for Cloud-Delivered Firewall Management Center, on page 247
- Troubleshoot LDAP authentication connections, on page 261
- Configure User Preferences, on page 261

CLI users

A CLI user is a user account that

- can be added as internal users or as external users on a LDAP or RADIUS server,
- is maintained separately on each managed device, and
- requires separate configuration for access to the Cloud-Delivered Firewall Management Center versus managed devices.

User account separation

When you add a user to the Cloud-Delivered Firewall Management Center, that user only has access to the Cloud-Delivered Firewall Management Center; you cannot then use that username to log directly into a managed device. You must separately add a user on the managed device.

Internal and external users

Managed devices support two categories of users—internal and external—for user access control.

Internal users authenticate through a local database on the device. External users authenticate through LDAP or RADIUS servers when they are not found in the local database. Both user types enable secure access management for network devices.

User authentication types

Managed devices support these user authentication methods:

- Internal user: The device checks a local database for user authentication.
- External user: If the user is not present in the local database, the system queries an external LDAP or RADIUS authentication server.

User roles

Web Interface User Roles

There are a variety of user roles in Security Cloud Control (Security Cloud Control): Read-Only, Edit-Only, Deploy-only, Admin, and Super Admin. User roles are configured for each user on each tenant. If a Security Cloud Control user has access to more than one tenant, they may have the same user ID but different roles on different tenants. A user may have a read-only role on one tenant and a Super Admin role on another. When the interface or the documentation refers to a Read-only user, Deploy Only, Edit Only, an Admin user, or a Super Admin user we are describing that user's permission level on a particular tenant. Note that you cannot create user roles in the Cloud-Delivered Firewall Management Center because it uses Security Cloud Control user roles.

Read Only

Read Only users can view all device configurations but cannot change them.

Deploy Only

Deploy Only users can audit queued changes made to device configurations and deploy them but cannot change them.

Edit Only

Edit Only users can make changes to all device configurations but cannot deploy them to devices.

Super Admin and Admin

Super Admin and Admin users can access everything in the product. As a Super Admin, you can create accounts for other users on a tenant and modify existing user roles. As an Admin, you cannot create accounts or modify user roles.

For more information about user roles in Security Cloud Control, refer to User Roles.

This table shows the mapping of user roles in On-Premises Firewall Management Center to their equivalent roles in the cloud-delivered Firewall Management Center in Security Cloud Control.

Tip We recommend that you read through the table only if you are familiar with the user roles in On-Premises Firewall Management Center.

Table 12: Secure Firewall Management Center and Cloud-Delivered Firewall Management Center User Role Mapping

On-Premises Firewall Management Center User Role	Equivalent Cloud-Delivered Firewall Management Center User Role	Capabilities
Access Admin, Discovery Admin, Intrusion Admin, Maintenance User	Edit Only	You can search, filter, or view these: • Access control policies and associated features • Intrusion policies • Intrusion rules • Network discovery rules • Custom detectors • Correlation policies • Objects • Rulesets • Interfaces • VPN configurations • Monitoring- and maintenance-related settings You can back up or restore a device but cannot deploy policies to the devices.
Administrator	Super Admin	You can access all features of the Cloud-Delivered Firewall Management Center and perform tasks, including create, read, modify, or delete policies or objects and deploy those changes to the devices. You can also edit user roles or create user records in Security Cloud Control.
Network Admin	Admin	You can access all features of the cloud-delivered Firewall Management Center and perform tasks, including create, read, modify, or delete policies or objects and deploy those changes to the devices. However, you cannot edit user roles or create user records in Security Cloud Control.

On-Premises Firewall Management Center User Role	Equivalent Cloud-Delivered Firewall Management Center User Role	Capabilities
Security Analyst, Security Analyst (Read Only)	Read Only	You can view device information, policies, objects, and their related settings but cannot do these: • Create or edit objects • Create or edit policies • Modify device configurations • Backup or restore devices
Security Approver	Deploy Only	You can view most settings and deploy staged changes to devices but cannot create or modify objects or policies.

Create a User Record with Your Security Cloud Control Username

Only a Security Cloud Control user with **Super Admin** privileges can create the Security Cloud Control user record. The **Super Admin** must create the user record with the same email address that was specified in the **Create Your Security Cloud Control Username** task above.

Use the following procedure to create a user record with an appropriate user role:

Procedure

Step 1 Login to Security Cloud Control.

Step 2 In the left pane, choose **Settings** > **User Management**.

Step 3 Click  to add a new user to your tenant.

Step 4 Provide the email address of the user.

Note
The user's email address must correspond to the email address of the Cisco Secure Log-On account.

Step 5 From the **Role** drop-down list, select the user's role.

Step 6 Click **OK**.

External authentication configuration for Cloud-Delivered Firewall Management Center

An external authentication configuration is a security setup that

- adds one or more external authentication objects,
- enables users to authenticate through external identity providers, and
- replaces local credential authentication with external authentication.

External authentication objects for Cloud-Delivered Firewall Management Center

An external authentication object is a security configuration element that

- enables the Cloud-Delivered Firewall Management Center to verify user credentials with LDAP or RADIUS servers,
- supports multiple objects for web interface access, allowing users from any configured object to authenticate, and
- restricts CLI access to a single external authentication object, permitting authentication only through the first object in the list.

External authentication objects can be used by the Cloud-Delivered Firewall Management Center and Firewall Threat Defense devices. You can either share the same object across different devices or create distinct objects for each device type.

For the Cloud-Delivered Firewall Management Center, enable the external authentication objects directly on the **System** (⚙) > **Users** > **External Authentication** tab. This setting affects only Cloud-Delivered Firewall Management Center usage. You do not need to enable it on this tab for managed device usage. For Firewall Threat Defense devices, enable the external authentication object in the platform settings that you deploy to the devices.

Define web interface users and CLI users separately in the external authentication object. For RADIUS, pre-configure the list of CLI usernames in the object. For LDAP, set up a filter on the LDAP server to match CLI users.



Note Users with CLI access can gain Linux shell access with the **expert** command. Linux shell users can obtain root privileges. This access creates a security risk. Make sure you:

- Restrict the list of users with CLI or Linux shell access.
- Do not create Linux shell users.

If you have five external authentication objects configured for web interface access, users from any of them can be authenticated to access the web interface. For CLI access, only the first external authentication object in the list is used for authentication.

You cannot use an LDAP object for CLI access if it is also configured for CAC authentication.

An external authentication object functions as a control point, permitting access only to users with valid credentials from specified sources.

LDAP

The Lightweight Directory Access Protocol (LDAP) is a directory service protocol that allows you to set up a directory on your network that organizes objects, such as user credentials, in a centralized location. Multiple applications can then access those credentials and the information used to describe them. If you ever need to change a user's credentials, you can change them in one place.

LDAP binding information

Microsoft has announced that Active Directory servers will start enforcing LDAP binding and LDAP signing in 2020. Microsoft is making these a requirement because when using default settings, an elevation of privilege vulnerability exists in Microsoft Windows that could allow a man-in-the-middle attacker to successfully forward an authentication request to a Windows LDAP server. For more information, refer to 2020 LDAP channel binding and LDAP signing requirement for Windows on the Microsoft support site.

If you have not done so already, we recommend you start using TLS/SSL encryption to authenticate with an Active Directory server.

RADIUS

RADIUS is a network authentication protocol that

- provides centralized authentication and authorization for network access,
- operates using client-server architecture for user credential verification, and
- supports accounting functions to track user sessions and network usage.

Remote Authentication Dial In User Service (RADIUS) is an authentication protocol used to authenticate, authorize, and account for user access to network resources. You can create an authentication object for any RADIUS server that conforms to RFC 2865.

Secure Firewall devices support the use of SecurID tokens. When you configure authentication by a server using SecurID, users authenticated against that server append the SecurID token to the end of their SecurID PIN and use that as their password when they log in. You do not need to configure anything extra on the Secure Firewall device to support SecurID.

By default, RADIUS uses port 1812 for authentication and port 1813 for accounting.

If you change the RADIUS authentication port, the RADIUS accounting port changes accordingly. Ensure that the Cloud-Delivered Firewall Management Center can connect to the RADIUS server on the new accounting port; otherwise, authentication delays may occur.

Although you cannot configure RADIUS accounting parameters, the Firewall Threat Defense uses port 1813 for RADIUS accounting on the same server used for authentication. If the RADIUS server is unreachable on port 1813, it can cause delays in logging in.

Add an LDAP external authentication object for the Security Cloud Control

Add an LDAP server to support external users for device management.

Before you begin

- You must specify DNS server(s) for domain name lookup on your device. Even if you specify an IP address for the LDAP server during this procedure, the server may return a URI for authentication that includes a hostname. A DNS lookup is required to resolve the hostname.
- If you are configuring an LDAP authentication object for use with CAC authentication, do not remove the CAC from your computer. You must have a CAC inserted at all times after enabling user certificates.

Procedure

Step 1 Choose a **Server Type** from the drop-down list.

Tip
If you click **Set Defaults** , the device populates the **User Name Template** , **UI Access Attribute** , **CLI Access Attribute** , **Group Member Attribute** , and **Group Member URL Attribute** fields with default values for the server type.

Step 2 For the **Primary Server**, enter a **Host Name/IP Address**. If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections. Now, change the **Port** from the default and enter the **Backup Server** parameters and **LDAP-Specific Parameters**.

a) Enter the **Base DN** for the LDAP directory you want to access. For example, to authenticate names in the Security organization at the Example company, enter `ou=security,dc=example,dc=com` . Alternatively click **Fetch DNs** , and choose the appropriate base distinguished name from the drop-down list.

b) (Optional) Enter the **Base Filter** . For example, if the user objects in a directory tree have a `physicalDeliveryOfficeName` attribute and users in the New York office have an attribute value of `NewYork` for that attribute, to retrieve only users in the New York office, enter `(physicalDeliveryOfficeName=NewYork)` .

If you are using CAC authentication, to filter only active user accounts (excluding the disabled user accounts), enter `(!(userAccountControl:1.2.840.113556.1.4.803:=2))` . This criteria retrieves user accounts within AD belonging to `ldpgrp` group and with `userAccountControl` attribute value that is not `2` (disabled).

c) Enter a **User Name** for a user who has sufficient credentials to browse the LDAP server. For example, if you are connecting to an OpenLDAP server where user objects have a `uid` attribute, and the object for the administrator in the Security division at your example company has a `uid` value of `NetworkAdmin` , you might enter `uid=NetworkAdmin,ou=security,dc=example,dc=com.`

d) Enter the user password in the **Password** and the **Confirm Password** fields.

e) (Optional) Click **Show Advanced Options** to configure these advanced options.

- **Encryption** —Click **None** , **TLS** , or **SSL** .

  If you change the encryption method after specifying a port, you reset the port to the default value for that method. For **None** or **TLS** , the port resets to the default value of 389

  - **None** — LDAP communication is unencrypted.

- **TLS** — LDAP communication is upgraded to TLS using the STARTTLS operation on default port 389
- **SSL** — LDAP communication is encrypted using TLS immediately upon connection (LDAPS) on default port 636.

- **SSL Certificate Upload Path** —For SSL or TLS encryption, click **Choose File** and choose the complete CA chain certificate.

  Note
  Do not choose a binary certificate (PKCS12, DER, and alike) file because Firewall Threat Defense does not support them.

  To remove the uploaded certificate, check the **Clear loaded certificate** check box. This option only appears when you have uploaded a certificate, and when you are in the Edit mode of the external authentication object.

  If you had previously uploaded a certificate and want to replace it, reupload the new certificate (complete CA chain), and redeploy the configuration to your devices to copy over the new certificate.

  Note
  TLS encryption requires a certificate on all platforms. We recommend that you **always** upload a certificate for SSL to prevent adversary-in-the-middle attacks.

- **User Name Template** —Provide a template that corresponds with your **UI Access Attribute** . For example, to authenticate all users who work in the Security organization of the Example company by connecting to an OpenLDAP server where the UI access attribute is `uid` , you might enter `uid=%s,ou=security,dc=example,dc=com` in the **User Name Template** field. For a Microsoft Active Directory server, you could enter `%s@security.example.com` .

  This field is required for CAC authentication.

- **Shell User Name Template** —Provide a template that corresponds with your **CLI Access Attribute** to authenticate CLI users. For example, to authenticate all users who work in the Security organization by connecting to an OpenLDAP server where the CLI access attribute is `sAMAccountName` , you might enter `%s` in the **Shell User Name Template** field.

- **Timeout (Seconds)** —Enter the number of seconds before rolling over to the backup connection, between 1 and 1024. The default is 30.

  Note
  The timeout range is different for Firewall Threat Defense and the Cloud-Delivered Firewall Management Center , so if you share an object, be sure not to exceed the Firewall Threat Defense 's smaller timeout range (1-30 seconds). If you set the timeout to a higher value, the Firewall Threat Defense LDAP configuration will not work.

Step 3 Configure **Attribute Mapping** to retrieve users based on an attribute.

- Enter a **UI Access Attribute** , or click **Fetch Attrs** to retrieve a list of available attributes. For example, on a Microsoft Active Directory Server, you may want to use the UI access attribute to retrieve users, because there may not be a `uid` attribute on Active Directory Server user objects. Instead, you can search the `userPrincipalName` attribute by typing `userPrincipalName` in the **UI Access Attribute** field.

- Set the **CLI Access Attribute** if you want to use a shell access attribute other than the user distinguished type. For example, on a Microsoft Active Directory Server, use the `sAMAccountName` CLI access attribute to retrieve CLI access users by typing `sAMAccountName` .

Step 4 (Optional) Configure **Group Controlled Access Roles**.

If you do not configure a user's privileges using group-controlled access roles, a user has only the privileges granted by default in the external authentication policy.

a) (Optional) In the fields that correspond to user roles, enter the distinguished name for the LDAP groups that contain users who should be assigned to those roles.

Any group you reference must exist on the LDAP server. You can reference static LDAP groups or dynamic LDAP groups. Static LDAP groups are groups where membership is determined by group object attributes that point to specific users, and dynamic LDAP groups are groups where membership is determined by creating an LDAP search that retrieves group users based on user object attributes. Group access rights for a role only affect users who are members of the group.

If you use a dynamic group, the LDAP query runs exactly as it is configured on the LDAP server. To prevent infinite loops caused by search syntax errors, the device limits recursions of a search to four.

Example:

Enter this in the **Administrator** field to authenticate names in the information technology organization at the Example company:

```
cn=itgroup,ou=groups, dc=example,dc=com
```

b) Choose a **Default User Role** for users that do not belong to any of the specified groups.

c) If you use static groups, enter a **Group Member Attribute** .

Example:

If the `member` attribute is used to indicate membership in the static group for default Security Analyst access, enter `member` .

d) If you use dynamic groups, enter a **Group Member URL Attribute** .

Example:

If the `memberURL` attribute contains the LDAP search that retrieves members for the dynamic group you specified for default Admin access, enter `memberURL` .

Step 5 (Optional) Set the **CLI Access Filter** to allow CLI users.

To prevent LDAP authentication of CLI access, leave this field blank. To specify CLI users, choose one of these methods:

- To use the same filter you specified when configuring authentication settings, check the check box of **Same as Base Filter** .
- To retrieve administrative user entries based on attribute value, enter the attribute name, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses. For example, if all network administrators have a `manager` attribute which has an attribute value of `shell` , you can set a base filter of `(manager=shell)` .

The usernames must be Linux-valid:

- Maximum 32 alphanumeric characters, plus period (.), hyphen (-), and underscore (_)

- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include at sign (@) or slash (/)

Note
Users with CLI access can gain Linux shell access with the **expert** command. "Linux shell users can obtain root privileges, which presents a security risk. Make sure that you restrict the list of users with CLI or Linux shell access.

Note
Do not create any internal users that have the same user name as users included in the **CLI Access Filter**. The only internal Cloud-Delivered Firewall Management Center user should be **admin** ; do not include an **admin** user in the **CLI Access Filter**.

Step 6 (Optional) Click **Test** to test connectivity to the LDAP server.

The test output lists valid and invalid user names. Valid user names are unique, and can include underscores (_), periods (.), hyphens (-), and alphanumeric characters. Note that testing the connection to servers with more than 1000 users only returns 1000 users because of UI page size limitations. If the test fails, refer to Troubleshoot LDAP authentication connections, on page 261 .

Step 7 (Optional) You can also enter **Additional Test Parameters** to test user credentials for a user who should be able to authenticate: enter a **User Name** `uid` and **Password** , and then click **Test**. Click **Save**. Enable use of this server. Refer to Enable external authentication for users on the CDO.

If you are connecting to a Microsoft Active Directory Server and supplied a UI access attribute in place of `uid` , use the value for that attribute as the user name. You can also specify a fully qualified distinguished name for the user.

Tip
If you mistype the name or password of the test user, the test fails even if the server configuration is correct. To verify that the server configuration is correct, click **Test** without entering user information in the **Additional Test Parameters** field first. If that succeeds, supply a user name and password to test with the specific user.

Example:

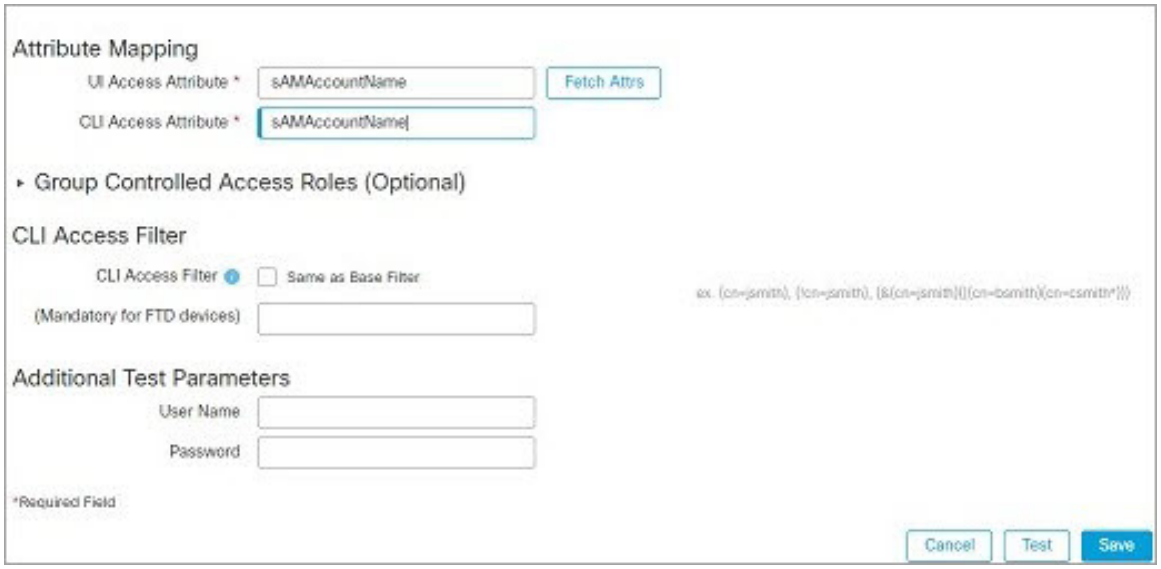
To test if you can retrieve the `JSmith` user credentials at the Example company, enter `JSmith` and the correct password.

Examples

Basic Example

This figures illustrate a basic configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in this example has an IP address of 10.11.3.4. The connection uses port 389 for access.

This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company.



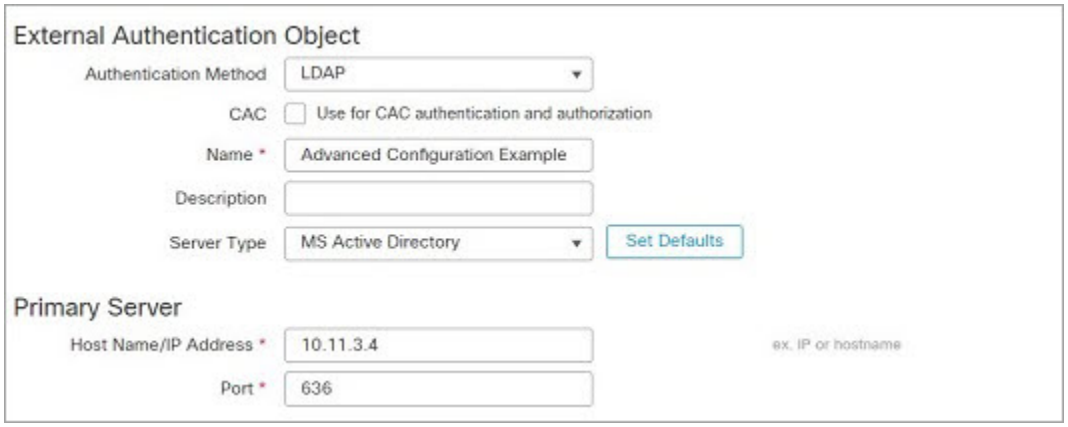
However, because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute. Choosing the MS Active Directory server type and clicking **Set Defaults** sets the UI Access Attribute to `sAMAccountName` . As a result, the system checks the `sAMAccountName` attribute for each object for matching user names when a user attempts to log into the system.

In addition, a **CLI Access Attribute** of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into a CLI account on the appliance.

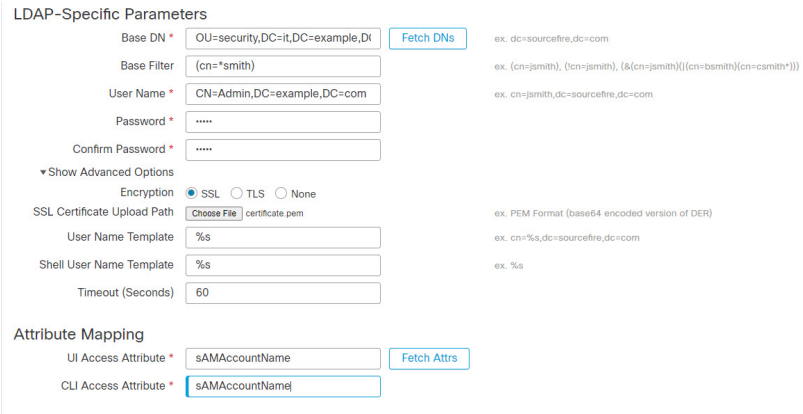
Note that because no base filter is applied to this server, the system checks attributes for all objects in the directory indicated by the base distinguished name. Connections to the server time out after the default time period (or the timeout period set on the LDAP server).

Advanced Example

This example illustrates an advanced configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in this example has an IP address of 10.11.3.4. The connection uses port 636 for access.



This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company. However, note that this server has a base filter of `(cn=*smith)` . The filter restricts the users retrieved from the server to those with a common name ending in `smith` .

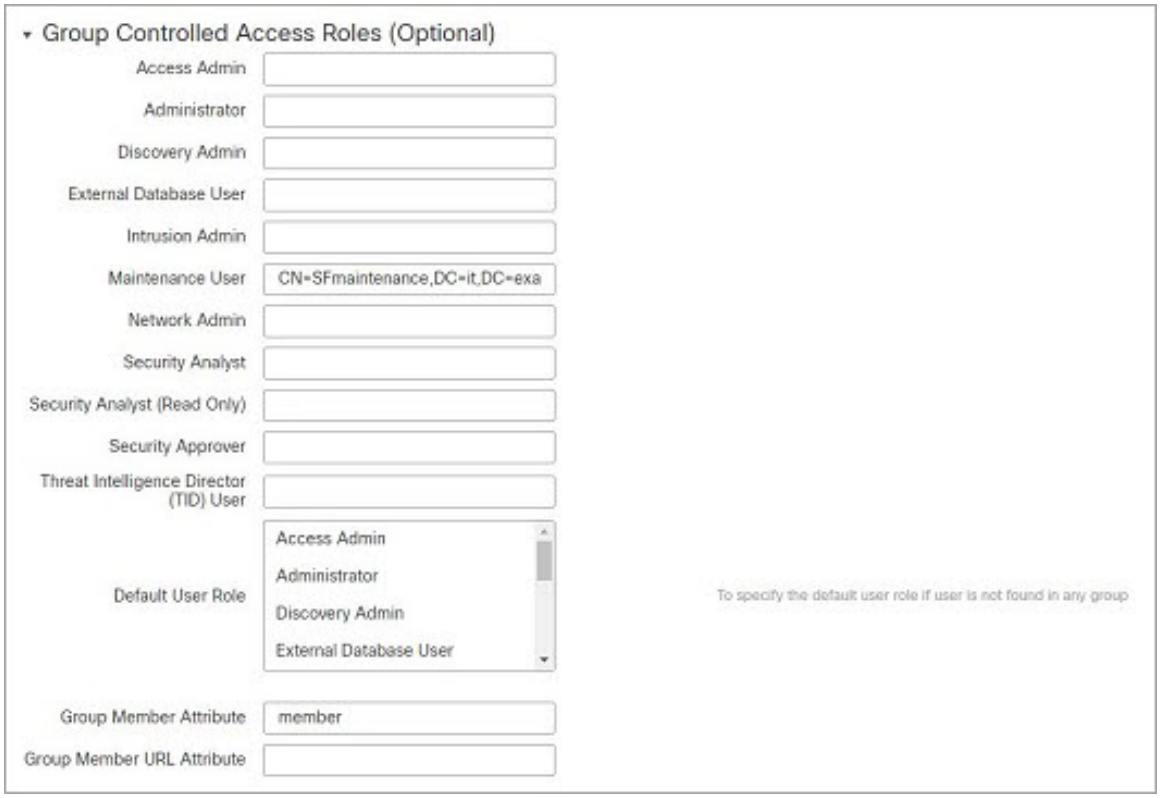


The connection to the server is encrypted using SSL and a certificate named `certificate.pem` is used for the connection. In addition, connections to the server time out after 60 seconds because of the **Timeout (Seconds)** setting.

Because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute. Note that the configuration includes a **UI Access Attribute** of `sAMAccountName` . As a result, the system checks the `sAMAccountName` attribute for each object for matching user names when a user attempts to log into the system.

In addition, a **CLI Access Attribute** of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into a CLI account on the appliance.

This example also has group settings in place. The Maintenance User role is automatically assigned to all members of the group with a `member` group attribute and the base domain name of `CN=SFmaintenance,DC=it,DC=example,DC=com` .



The **CLI Access Filter** is set to be the same as the base filter, so the same users can access the appliance through the CLI as through the web interface.



Add a RADIUS external authentication object for Security Cloud Control

Add a RADIUS server to support external users for device management.

Note that supports only these six attributes in the access request of the RADIUS authentication:

- `NAS-IP-Address` : Used in UI and CLI authentication
- `NAS-Identifier` : Used in CLI authentication
- `NAS-Port` : Used in UI and CLI authentication

- `NAS-Port-Type` : Used in CLI authentication
- `Service-Type` : Used in UI and CLI authentication
- `Calling-Station-Id` : Used in CLI authentication

Note

- `NAS-Port` attribute: The attribute value is dynamically generated and cannot be configured.
- `NAS-Port-Type` attribute: The server always sets the attribute value to `Virtual`, regardless of whether the device is virtual or physical. of whether the is virtual or physical.
- `NAS-Identifier` attribute: The attribute value is always `sshd`.
- `Service-Type` attribute: supports only `Login` and `Authenticate Only` service types

Procedure

Step 1 Choose **System** (⚙) > **Users** and click **External Authentication**.

Step 2 Click **Add icon** (+). In the **Add External Authentication Object** window, set the **Authentication Method** to **RADIUS**, and enter a unique **Name** and optional **Description**.

Step 3 Check the **RADIUS Server-Enabled Message Authenticator** check box. This requires the Message-Authenticator attribute in all RADIUS responses and ensures that every response from the RADIUS server is securely verified by the Firewall Threat Defense.

The feature is enabled by default for new RADIUS servers. Enable it for existing servers after the upgrade. Enable message authenticators to protect your firewalls from potential attacks. Ensure that your RADIUS server has the Message-Authenticator configuration.

Note
The Message-Authenticator attribute persists whether the feature is enabled or disabled. The server always requires this attribute. When the check box is enabled in the configuration, it also becomes a requirement from the Cloud-Delivered Firewall Management Center.

Step 4 For the **Primary Server**, enter a **Host Name/IP Address**. Then, change the **Port** from the default and then enter the **RADIUS Secret Key**.

Step 5 (Optional) Enter the **Backup Server** parameters, then provide the **RADIUS-Specific Parameters**.

a) Enter the **Timeout** in seconds before retrying the primary server, between 1 and 1024. The default is 30.
b) Enter the **Retries** before rolling over to the backup server. The default is 3.
c) In the fields that correspond to user roles, enter the name of each user or identifying attribute-value pair that should be assigned to those roles.

Separate usernames and attribute-value pairs with commas.

Example:

If you know all users who should be Security Analysts have the value `Analyst` for their `User-Category` attribute, you can enter `User-Category=Analyst` in the **Security Analyst** field to grant that role to those users.

Example:

To grant the Administrator role to the users `jsmith` and `jdoe`, enter `jsmith, jdoe` in the **Administrator** field.

Example:

To grant the Maintenance User role to all users with a `User-Category` value of `Maintenance`, enter `User-Category=Maintenance` in the **Maintenance User** field.

d) Select the **Default User Role** for users that do not belong to any of the specified groups.

After you change a user's role, save and deploy the updated external authentication object. Remove the user from the **Users** screen. The system adds the user when they log in again.

Step 6 (Optional) **Define Custom RADIUS Attributes**.

If your RADIUS server returns values for attributes not included in the `dictionary` file in `/etc/radiusclient/`, define those attributes. You must define them if you plan to use the attributes to set roles for users.You can locate the attributes returned for a user by looking at the user's profile on your RADIUS server.

Note
Before you configure custom RADIUS attributes, obtain the Attribute ID and Type from the RADIUS server administrator. If there is a mismatch, the RADIUS packet capture displays the custom RADIUS attribute as `Unassigned Attribute`.

a) Enter an **Attribute Name**.

Enter an attribute name using alphanumeric characters. Separate words in an attribute name with dashes instead of spaces.

b) Enter the **Attribute ID** as an integer.

Enter an integer for the attribute ID. Make sure the attribute ID does not conflict with any existing IDs in the `etc/radiusclient/dictionary` file.

c) Choose the **Attribute Type** from the drop-down list.

You also specify the type of attribute: string, IP address, integer, or date.

d) Click **Add** to add the custom attribute.

When you create a RADIUS authentication object, a new dictionary file for that object is created on the device in the `/var/sf/userauth` directory. Any custom attributes you add are added to the dictionary file.

Example:

If a RADIUS server is used on a network with a Cisco router, you might want to use the `Ascend-Assign-IP-Pool` attribute to grant a specific role to all users logging in from a specific IP address pool. `Ascend-Assign-IP-Pool` is an integer attribute that defines the address pool where the user is allowed to log in, with the integer indicating the number of the assigned IP address pool.

To declare that custom attribute, you create a custom attribute with an attribute name of `Ascend-IP-Pool-Definition`, an attribute ID of `218`, and an attribute type of `integer`.

You could then enter `Ascend-Assign-IP-Pool=2` in the **Security Analyst (Read Only)** field to grant read-only security analyst rights to all users with an `Ascend-IP-Pool-Definition` attribute value of `2`.

Step 7 (Optional) In the **CLI Access Filter** area **Administrator CLI Access User List** field, enter the usernames that should have CLI access, separated by commas.

Enter usernames that match those on the RADIUS server. Usernames must be valid Linux usernames:

- Maximum of 32 alphanumeric characters, plus period (.), hyphen (-), and underscore (_).

  (_)

- All letters must be lowercase.
- "The username cannot start with a hyphen (-), cannot be composed entirely of numbers, and cannot include an at sign (@) or a slash (/).

Leave the field blank to disable RADIUS authentication for CLI access.

Note
Users with CLI access can gain Linux shell access with the **expert** command. Linux shell users can obtain root privileges, which can present a security risk. Make sure that you restrict the list of users with CLI or Linux shell access.

Note
Remove any internal users that have the same user name as users included in the shell access filter. For the Cloud-Delivered Firewall Management Center, the only internal CLI user is **admin**, so do not also create an **admin** external user.

Step 8 (Optional) Click **Test** to test Cloud-Delivered Firewall Management Center connectivity to the RADIUS server.

Step 9 (Optional) You can also enter **Additional Test Parameters** to test user credentials for a user who should be able to authenticate: enter a **User Name** and **Password**, and then click **Test**. Click **Save**. For more information, refer to Enable external authentication for users on the CDO.

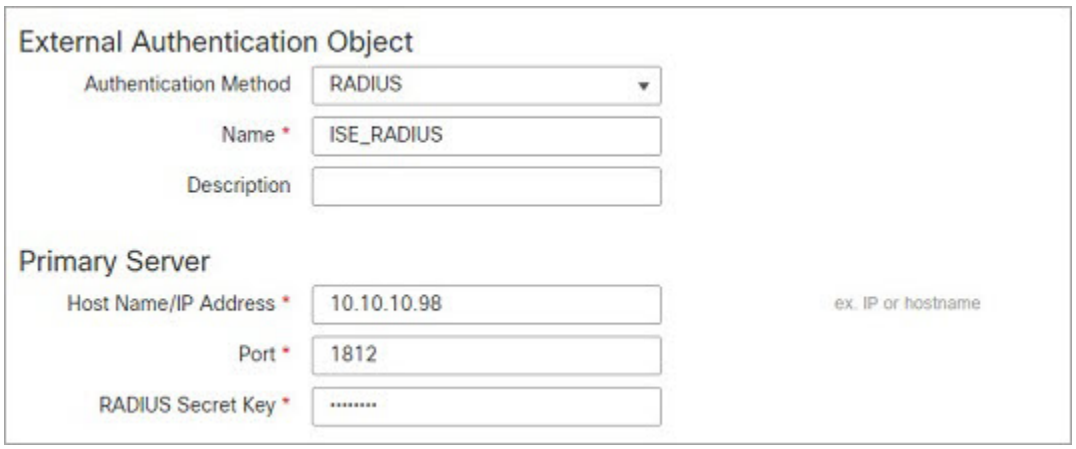
Tip
If you mistype the name or password of the test user, the test fails even if the server configuration is correct. To verify that the server configuration is correct, click **Test** without entering user information in the **Additional Test Parameters** field first. If that succeeds, supply a user name and password to test with the specific user.

Example:

To test if you can retrieve the `JSmith` user credentials at the Example company, enter `JSmith` and the correct password.

Simple user role assignments

This figure illustrates a sample RADIUS login authentication object for a server running Cisco Identity Services Engine (ISE) with an IP address of 10.10.10.98 on port 1812. No backup server is defined.



This example shows RADIUS-specific parameters, including the timeout (30 seconds) and number of failed retries before the Secure Firewall System attempts to contact the backup server, if any.

This example illustrates important aspects of RADIUS user role configuration:

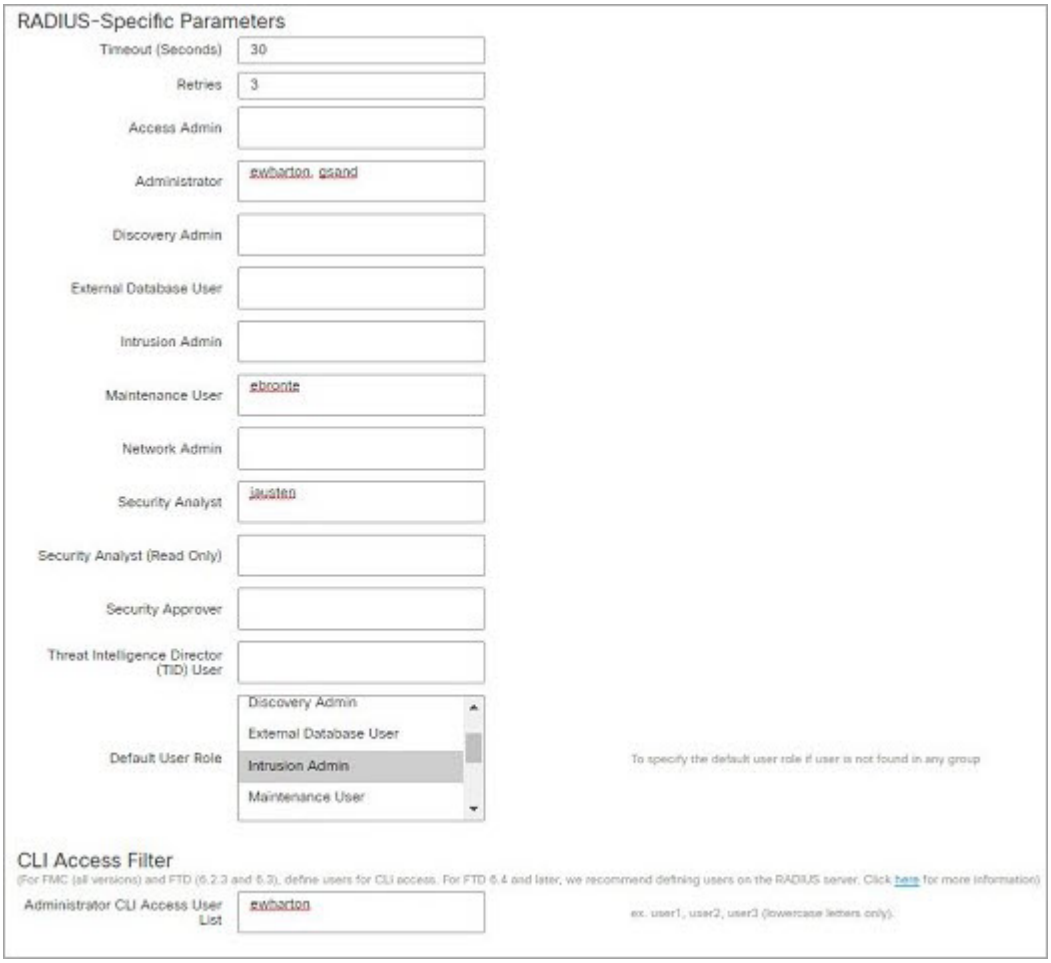
Users `ewharton` and `gsand` are granted web interface Administrative access.

The user `cbronte` is granted web interface Maintenance User access.

The user `jausten` is granted web interface Security Analyst access.

The user `ewharton` can log into the device using a CLI account.

This graphic depicts the role configuration for the example:

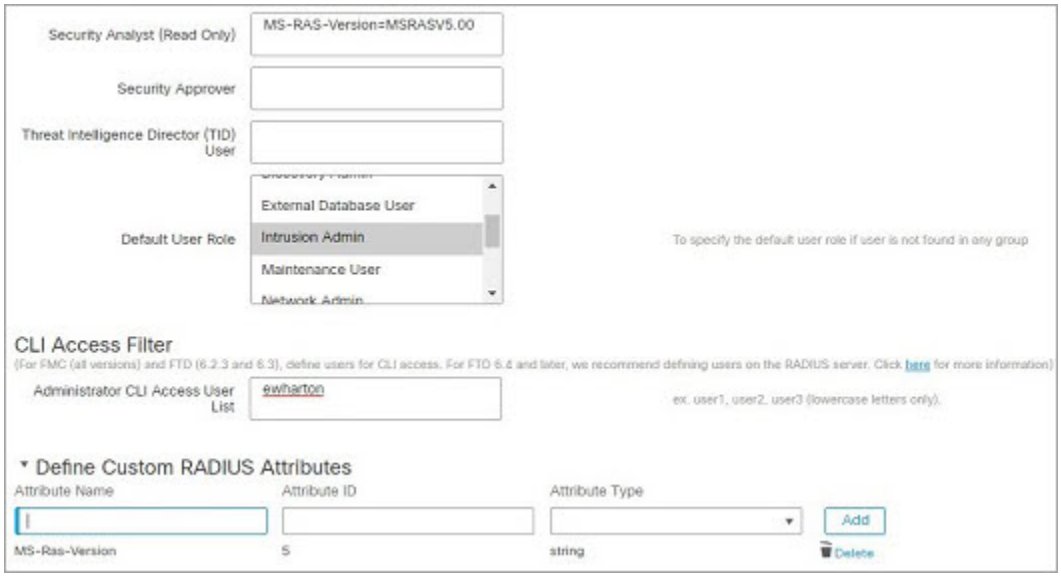


Roles for users matching an attribute-value pair

You can use an attribute-value pair to identify users who should receive a particular user role. If the attribute you use is a custom attribute, you must define the custom attribute.

This figure illustrates the role configuration and custom attribute definition in a sample RADIUS login authentication object for the same ISE server as in the previous example.

In this example, however, the `MS-RAS-Version` custom attribute is returned for one or more of the users because a Microsoft remote access server is in use. Note the `MS-RAS-Version` custom attribute is a string. In this example, all users logging in to RADIUS through a Microsoft version 5.00 remote access server should receive the Security Analyst (Read Only) role, so you enter the attribute-value pair of `MS-RAS-Version=MSRASV5.00` in the **Security Analyst (Read Only)** field.



Troubleshoot LDAP authentication connections

This reference explains how to troubleshoot LDAP authentication connection issues, identify common failure causes, and apply solutions to restore communication between Cisco devices and LDAP servers.

Configure User Preferences

Depending on your user role, you can specify certain preferences for your user account.

Change the web interface appearance

Change the web interface appearance by choosing from available themes or enable new design options for the light theme variant.

You can change the way the web interface appears when you are using the light theme. Theme options and design toggles may vary based on your system version and available features.

Before you begin

Follow these steps to change the web interface appearance:

Procedure

Step 1 From the drop-down list under your username, choose **User Preference**.

Step 2 Under **Light theme variant**, click the **Use new design** toggle button to enable or disable the new design.

Enable the toggle to use the left navigation menu and disable it to use the top navigation menu.

Note
This toggle works only if you are using the light theme. If you want to switch between dark theme and light theme, you must do this from the **General Preference** page in Security Cloud Control. For more information, refer to Change the Security Cloud Control Web Interface Appearance.

The web interface appearance changes according to your selected theme or design preference, providing a personalized and comfortable user experience.

Set your default time zone

Set your default time zone so that all times shown in the web interface for your user account reflect your local time preferences. This helps you view scheduled tasks and dashboards in your preferred time zone.

This setting determines the times displayed in the web interface for your user account, such as for task scheduling and dashboard viewing. It does not change the system time, affect other users, or modify data stored in the system, which generally uses UTC.

- The Time Zone function in User Preferences assumes that the system clock is set to UTC. Do not change the system time. Changing the system time from UTC is not supported and will require reimaging the device to recover from an unsupported state.
- This feature does not affect the time zone used for time-based policy application. Set the time zone for a device in **Devices** > **Platform Settings**.

Before you begin

Follow these steps to set your default time zone:

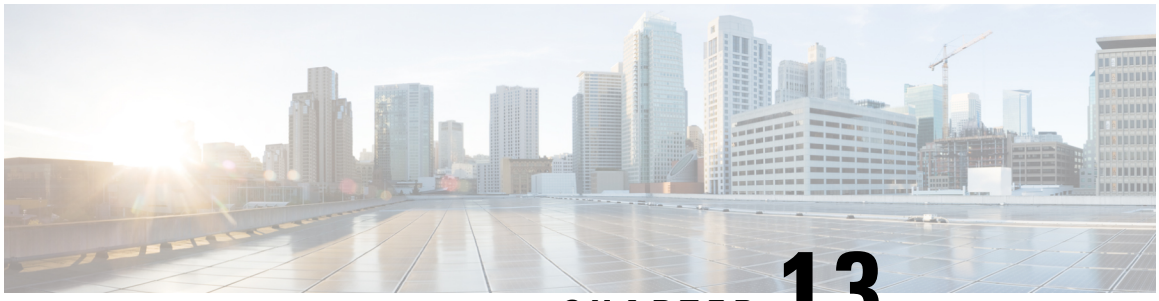
Procedure

Step 1 From the drop-down list under your user name, choose **User Preferences**.

Step 2 Click the **TIME Zone** drop-down list.

Step 3 Choose the continent or country and the state name that corresponds with the time zone you want to use.

After completing these steps, your user account will display times in the web interface according to your selected time zone. System time and other users are not affected.



CHAPTER 13

Updates

- Content updates, on page 263
- Product upgrades, on page 264
- Guidelines for content updates, on page 264
- Vulnerability database (VDB) updates, on page 264
- Geolocation database (GeoDB) updates, on page 266
- Intrusion rule updates, on page 268

Content updates

Content updates keep your deployment current with changing security and contextual information. The system uses this information to evaluate network activity and respond to changes in the threat landscape. Automatic content updates are often enabled by initial configuration or when you enable the related feature.

Types of content updates

Cloud-Delivered Firewall Management Center uses these types of content updates:

- Vulnerability database (VDB): A database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The system uses the VDB to help determine whether a particular host increases your risk of compromise.

  Refer to Vulnerability database (VDB) updates, on page 264.

- Geolocation database (GeoDB): Maps IP addresses to countries and continents.

  Refer to Geolocation database (GeoDB) updates, on page 266.

- Intrusion rules (LSP/SRU): Intrusion rule updates provide new and updated intrusion and preprocessor rules and can modify rule states, categories, variables, and default policy settings.

  Refer to Intrusion rule updates, on page 268.

- Security Intelligence feeds: Collections of IP addresses, domain names, and URLs used to filter matching traffic.

  Refer to List and feed updates for security intelligence feeds, on page 3025.

- URL categories and reputations: Classification and risk information used to control access to websites.

  Refer to Enable URL filtering using category and reputation, on page 2096.

Product upgrades

This guide does not include procedures for upgrading firewall software or chassis. Instead, refer to the Secure Firewall Threat Defense upgrade guides for Cloud-Delivered Firewall Management Center.

Guidelines for content updates

These guidelines apply to all content updates. Refer to the topics for each update type for update-specific guidance.

Release information

Read any release notes or advisory text that accompanies a content update. This information describes compatibility, prerequisites, new capabilities, behavior changes, and warnings.

When to perform content updates

Follow these guidelines when performing and scheduling content updates:

- Scheduled updates: Review scheduled updates to make sure they run at the right time for your environment.
- Traffic inspection and flow: Perform updates that might interrupt traffic during maintenance windows.

Content updates are scheduled in Coordinated Universal Time (UTC), which remains constant year-round. Scheduled updates do not adjust automatically for local time changes, such as daylight saving time or summer time. For example, an update scheduled for 2:00 a.m. during standard time runs at 3:00 a.m. after the local clock moves forward.

Vulnerability database (VDB) updates

The VDB is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The system uses the VDB to help determine whether a particular host increases your risk of compromise.

VDB versions

We periodically release VDB updates. Each update replaces the previous version.

Starting with VDB 357, you can install an earlier VDB as far back as the baseline VDB for Cloud-Delivered Firewall Management Center.

VDB update duration

The time required to update the VDB and its associated mappings on the Cloud-Delivered Firewall Management Center depends on the number of hosts in your network map. Allow approximately one minute for every 1000 hosts.

Scheduled VDB updates

Initial configuration automatically downloads and installs the latest VDB as a one-time operation and creates a weekly task to download future updates. Review the schedule and adjust as needed. To update the VDB and deploy configurations automatically, create separate tasks. For more information, refer to Scheduling, on page 471.

Deployment after VDB updates

Deploy configuration changes after a VDB update for updated application detectors and operating system fingerprints to take effect. Updated vulnerability information takes effect without deploy.

Caution The first deploy after a VDB update usually restarts Snort because updated application detectors and operating system fingerprints require a restart. Restarting Snort briefly interrupts traffic flow and inspection on all devices, including those configured for high availability or scalability. Interface configurations determine whether traffic drops or passes without inspection during the interruption. When you deploy without restarting Snort, resource demands may result in a small number of packets dropping without inspection.

VDB release information

For VDB 343 and later, Cisco Secure Firewall Application Detectors provides searchable application detector information. The release notes describe changes in each VDB release.

VDB lite

For the Secure Firewall 220, the system installs a smaller VDB (also called VDB lite). This smaller VDB contains the same applications, but fewer detection patterns. Devices using the smaller VDB can miss some application identification versus devices using the full VDB.

Manually update the VDB

Perform an on-demand update to keep the VDB current or to install an older VDB.

Before you begin

If you are installing an older VDB, download it manually from https://www.cisco.com/go/firepower-software. Choose any Cloud-Delivered Firewall Management Center model, then open the Coverage and Content Updates page.

Starting with VDB 357, you can install any VDB as far back as the baseline VDB for the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **VDB Updates**.

Step 2 Choose how you want to get the VDB onto the Cloud-Delivered Firewall Management Center.

- Direct download: Click **Download Updates**.

- Manual upload: Click **Upload Update**, click **Choose File**, select the VDB, and click **Upload**.

Step 3 Install the VDB.

a) Next to the VDB update you want to install, click either the **Install** icon (for a newer VDB) or the **Rollback** icon (for an older VDB).
b) Choose the Cloud-Delivered Firewall Management Center.
c) Click **Install**.

Monitor update progress in the Message Center. Do not use mapped-vulnerability features while the VDB is updating. If the Message Center shows no progress for several minutes or reports that the update failed, do not restart it. Contact Cisco TAC.

Step 4 Verify update success.

The current version is displayed on the VDB update page.

The system uses the new vulnerability information. However, you must deploy before updated application detectors and operating system fingerprints can take effect.

What to do next

- Deploy configuration changes.
- If configurations reference vulnerabilities, application detectors, or fingerprints that are no longer available, review those configurations to confirm that they handle traffic as expected. A scheduled VDB update can undo a rollback. To retain the older VDB, change the scheduled task or delete any newer VDB packages.

Geolocation database (GeoDB) updates

The geolocation database (GeoDB) maps IP addresses to geographic locations so that you can view and filter traffic by location.

GeoDB versions

We periodically release GeoDB updates. Each update replaces the previous version.

Scheduled GeoDB updates

Initial configuration schedules weekly updates to keep geolocation information accurate. Review the schedule and adjust it as needed. Refer to Schedule GeoDB updates, on page 267.

Deployment after GeoDB updates

The Cloud-Delivered Firewall Management Center automatically updates the GeoDB on managed devices, so deploy is usually unnecessary. If an update adds a new country, which is rare, deploy configuration changes as soon as possible. This allows the new country to count as part of its continent. For example, if an update adds Country to Continent, rules that filter based on "Continent" do not match traffic through Country until you deploy.

Schedule GeoDB updates

Schedule recurring GeoDB updates to keep the installed geolocation information current.

Initial configuration creates a weekly schedule for GeoDB updates. Review the schedule and adjust the start time as needed.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Geolocation Updates**.

Step 2 Configure recurring GeoDB updates.

a) Under **Recurring Geolocation Updates**, check **Enable Recurring Weekly Updates from the Support Site**.

b) Specify the **Update Start Time**.

Step 3 Click **Save**.

Manually update the GeoDB

Perform an on-demand update to keep the GeoDB current.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Geolocation Updates**.

Step 2 Under **One-Time Geolocation Update**, choose how you want to get the GeoDB onto the Cloud-Delivered Firewall Management Center.

- Direct download: Choose **Download and install geolocation update from the Support Site**.
- Manual upload: Choose **Upload and install geolocation update**, click **Choose File**, and select the GeoDB.

Step 3 Click **Import**.

Monitor update progress in the Message Center.

Step 4 Verify update success.

The current version is displayed on the GeoDB update page.

The system uses the new geolocation information. However, you must deploy before new countries can take effect. Until you deploy, the new country does not count as part of its continent. For example, if an update adds Country to Continent, rules that filter based on "Continent" do not match traffic through Country until you deploy.

What to do next

If the update adds a new country, deploy configuration changes.

Intrusion rule updates

An intrusion rule update is a cumulative package that provides new and modified intrusion and preprocessor rules, and modifies the policies that use them.

Intrusion rule update versions

As new vulnerabilities become known, the Talos Intelligence Group releases intrusion rule updates. Keep intrusion rules up to date. Each update replaces the previous version, and you can import only an update that is newer than the currently installed version.

Contents of intrusion rule updates

An intrusion rule update can contain these types of content:

- Rules and rule states: An update can add, modify, or delete intrusion and preprocessor rules. The default state of a new rule can differ between system-provided intrusion policies. For example, a new rule can be enabled in the Security over Connectivity policy and disabled in the Connectivity over Security policy. An update can also change the default state of an existing rule.
- Rule categories: An update can contain new rule categories. Importing the update adds all new categories.
- Preprocessor and advanced settings: An update can change advanced settings in system-provided intrusion policies and preprocessor settings in system-provided network analysis policies. It can also change the default values of advanced preprocessing and performance options in access control policies.
- Variables: An update can change the default values of existing variables without overriding customized values. It also adds new variables.

Policy changes from intrusion rule updates

Intrusion rule updates can affect system-provided and custom network analysis and intrusion policies, as well as the access control policies that use them. The effects depend on the policy type.

- System-provided policies: Changes to system-provided network analysis and intrusion policies and advanced access control settings take effect after you redeploy the configuration.
- Custom policies: Every custom network analysis and intrusion policy is based directly or indirectly on a system-provided policy. For each custom policy, you can prevent rule updates from automatically applying changes from its base policy. You can then apply the base-policy changes manually on a schedule independent of rule update imports. Updates do not override customized settings.



Caution Importing an intrusion rule update discards all cached changes to network analysis and intrusion policies. Before you import an update, use the Rule Updates page to identify policies with cached changes and the users who made them.

Scheduled intrusion rule updates

Initial configuration creates a daily schedule for intrusion rule updates. Review the schedule and adjust the frequency as needed. Refer to Schedule intrusion rule updates, on page 269.

Deployment after intrusion rule updates

Deploy configuration changes after an intrusion rule update for the update's changes to take effect. You can configure an import to deploy automatically to affected devices. Automatic deployment is especially useful when you allow updates to modify system-provided base intrusion policies.

Caution Although a rule update by itself does not restart Snort when you deploy, other changes you have made may. Restarting Snort briefly interrupts traffic flow and inspection on all devices, including those configured for high availability or scalability. Interface configurations determine whether traffic drops or passes without inspection during the interruption. When you deploy without restarting Snort, resource demands may result in a small number of packets dropping without inspection.

Schedule intrusion rule updates

Schedule recurring intrusion rule updates to keep the installed intrusion rules current.

Before you begin

Make sure your process for updating intrusion rules complies with your security policies.

Initial configuration creates a daily schedule for intrusion rule updates. Review the schedule and adjust the frequency as needed. You can also configure automatic deployment after each rule update. Automatic deployment can disrupt traffic inspection and flow.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Rule Updates**.

Step 2 Configure recurring rule updates.

a) Under **Recurring Rule Update Imports**, check **Enable Recurring Rule Update Imports**.
b) Specify the **Import Frequency** and start time.
c) (Optional) Check **Deploy all policies to targeted devices after rule update completes**.

Step 3 Click **Save**.

Manually update intrusion rules

Perform an on-demand update to keep intrusion rules current.

Before you begin

Make sure your process for updating intrusion rules complies with your security policies.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Rule Updates**.

Step 2 Configure the one-time rule update.

a) Under **One-Time Rule Update/Rules Import**, choose how to update intrusion rules.

- Direct download: Choose **Download new rule update...**.
- Manual upload: Choose **Rule update or text rule file...**, click **Choose File**, and select the intrusion rule update.

b) (Optional) Check **Reapply all policies after the rule update import completes**.

Step 3 Click **Import**.

Monitor update progress in the Message Center. Even if the Message Center shows no progress for several minutes or indicates that the update has failed, do not restart the update. Instead, contact Cisco TAC.

Step 4 Verify update success.

The current version is displayed on the rule update page.

What to do next

If you did not deploy configuration changes as part of the update, deploy them now.

Guidelines for importing custom Snort 2 rules

Use these guidelines to import custom standard text rules for use in Snort 2 intrusion policies.

File and rule syntax requirements

Import rules using a rule file that meets these requirements:

- Use a plain text file encoded in ASCII or UTF-8.
- Use only alphanumeric characters, spaces, underscores (_), periods (.), and hyphens (-) in the file name.
- To import a rule in the deleted state, precede it with one number sign (#). A rule preceded by two number signs (##) is not imported.
- Do not use escape characters in rules.
- Limit each source or destination port list to 64 characters. A longer list causes the import to fail.

For more information on writing custom rules, refer to *Custom Snort 2 Intrusion Policies for Access Control* and to the Snort manual at http://www.snort.org.

SID, GID, and revision number requirements

Follow these requirements for SIDs, GIDs, and revision numbers:

- Do not specify a GID. A standard text rule receives GID 1.
- Leave the SID and revision number unspecified when importing a rule for the first time. The rule receives an available SID of 1000000 or greater and a revision number of 1.
- If you must specify a SID for a new rule, use an unassigned value from 1000000 through 2147483647 (the 32-bit limit).
- When importing an updated or previously deleted rule, include its assigned SID and a revision number greater than its current revision number.

To determine the current revision number, edit the imported rule. Deleting a custom rule moves it to the deleted rule category and increments its revision number. You can then import a later revision to reinstate the rule.

Import custom Snort 2 rules

Import custom standard text rules so you can enable them in Snort 2 intrusion policies.

This procedure applies only to custom standard text rules for Snort 2 intrusion policies. For Snort 3, refer to the rule-tuning topics in *Custom Snort 3 Intrusion Policies for Access Control*.

Before you begin

- Make sure your custom rule file follows the guidelines in Guidelines for importing custom Snort 2 rules, on page 270.
- Make sure your process for importing custom rules complies with your security policies.

Procedure

Step 1 Choose **System** (⚙) > **Content Updates** > **Rule Updates**.

You can also click **Import Rules** in the intrusion rules editor (**Objects** > **Intrusion Rules**).

Step 2 (Optional) Delete all existing custom rules.

To replace all existing custom rules with the rules in the new file, click **Delete All Local Rules**. The deleted rules move to the deleted rule category, and their revision numbers increase.

Step 3 Import the rule file.

a) Under **One-Time Rule Update/Rules Import**, choose **Rule update or text rule file to upload and install**.
b) Click **Choose File**, select your rule file, and click **Import**.

You can monitor import progress in the Message Center. Even if the Message Center shows no progress for several minutes or indicates that the import has failed, do not restart the import. Instead, contact Cisco TAC.

- Imported rules are added to the local rule category in a disabled state.
- Imported rules receive Generator ID (GID) 1.

- New imported rules receive an available Snort ID (SID) of 1000000 or greater and a revision of 1.

What to do next

- Edit Snort 2 intrusion policies and enable the rules you imported.

  Do not enable an imported rule that uses the deprecated `threshold` keyword in an intrusion policy that also uses intrusion event thresholding. This combination causes policy validation to fail.

- Deploy configuration changes.

Intrusion rule update logs

The Rule Update Log on the Rule Updates page records each intrusion rule update and custom Snort 2 rule import, including its time, user, and status. Each log also identifies the affected rules and components. Deleting a log does not delete the imported objects.

Intrusion rule update log fields

This table explains the fields in intrusion rule update logs.

Table 13: Intrusion rule update log fields

Field	Description
Action	Identifies how the import affected the object.
Default action	For a rule, the default action defined by the update is Pass, Alert, or Drop. The field is blank for other object types.
Details	Identifies the affected component or rule. For a changed rule, the field displays the GID, SID, and previous revision number in `GID:SID:Rev` format. The field is blank for a rule that has not changed.
GID	The generator ID for a rule. For example: • A GID of 1 identifies a standard text rule. • A GID of 3 identifies a shared object rule.
Name	The name of the imported object. For a rule, the name corresponds to the rule's Message field. For a rule update component, the field displays the component name.
Policy	For an imported rule, the field displays `All` when the import succeeds and the rule can be enabled in all applicable system-provided default intrusion policies. The field is blank for other object types.
Rev	The revision number for a rule.
Rule update	The rule update file name.
SID	The Snort ID for a rule.
Time	The time and date the import began.

Field	Description
Type	Identifies the imported object as a rule update component, rule, or policy apply. A policy apply record indicates that the option to reapply all policies after the import was enabled.

Action values

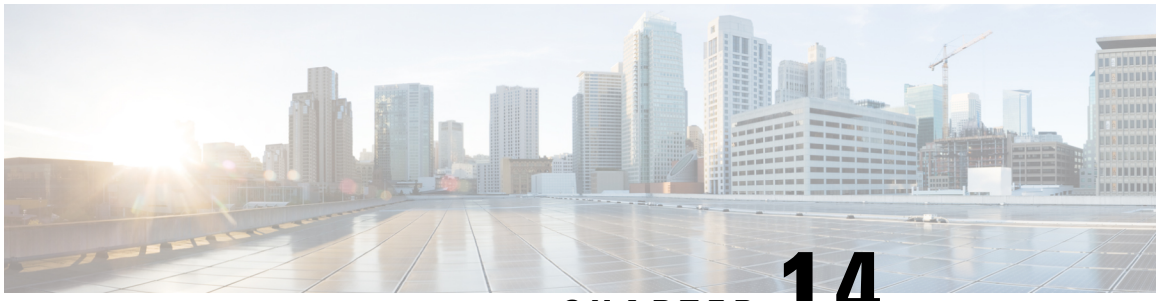
The Action field in intrusion rule update logs uses these values.

Table 14: Action values in intrusion rule update logs

Action	Meaning
new	The rule was stored for the first time.
changed	A rule update component was modified, or a rule was imported with a higher revision number and the same GID and SID.
collision	The import was skipped because the revision conflicts with an existing component or rule.
deleted	The rule was deleted from the rule update.
enabled	A rule, preprocessor, or other feature was enabled in a system-provided default policy.
disabled	A rule was disabled in a system-provided default policy.
drop	A rule was set to Drop and Generate Events in a system-provided default policy.
error	The update or import failed.
apply	The option to reapply all policies after the import was enabled.

Searching intrusion rule update logs

A search from the detailed log view searches the entire database, not only the selected import. Use the time range to limit the results to relevant records.



CHAPTER 14

Licenses

This chapter provides in-depth information about the different license types, service subscriptions, licensing requirements and more.

- Licenses, on page 275
- Prerequisites for licensing, on page 289
- Create a Cisco account, on page 290
- Create a smart account and add licenses, on page 291
- Configure Smart Licensing, on page 292
- Configure Legacy Cloud-Delivered Firewall Management Center PAK-Based Licenses, on page 300
- Additional information about licensing, on page 301
- Overview, on page 301
- Requirements, on page 302
- Supported Platforms and Versions, on page 302
- How to Implement Universal Permanent Licensing, on page 302
- Universal Permanent License Status, on page 308
- Maintain Your UPLR Deployment, on page 309
- Deactivate and Return a Universal Permanent License, on page 309
- Troubleshooting, on page 312
- For Assistance, on page 313

Licenses

Cisco Smart Licensing is a flexible licensing model that provides you with an easier, faster, and more consistent way to purchase and manage software across the Cisco portfolio and across your organization. And it's secure—you control what users can access. With Smart Licensing you get:

- **Easy Activation**: Smart Licensing establishes a pool of software licenses that can be used across the entire organization—no more PAKs (Product Activation Keys).
- **Unified Management**: Cisco License Central provides a complete view into all of your Cisco products and services in an easy-to-use portal, so you always know what you have and what you are using.
- **License Flexibility**: Your software is not node-locked to your hardware, so you can easily use and transfer licenses as needed.

To use Smart Licensing, you must first set up a Smart Account on Cisco License Central.

For a more detailed overview on Cisco Licensing, visit the Licensing Hub.

Smart Software Manager and accounts

When you purchase one or more licenses, you manage them in the Smart Software Manager: https://software.cisco.com/#module/SmartLicensing. The Smart Software Manager lets you create a primary account for your organization. If you do not yet have an account, click the link to set up a new account. For instructions, refer to Create a Cisco Account.

By default, your licenses are assigned to the Default Virtual Account under your primary account. As the account administrator, you can create additional virtual accounts; for example, for regions, departments, or subsidiaries. Multiple virtual accounts help you manage large numbers of licenses and devices.

You manage licenses by virtual account. Only that virtual account's devices can use the licenses assigned to the account. If you need additional licenses, you can transfer an unused license from another virtual account. You can also transfer devices between virtual accounts.

Licensing for Management Center and devices

The Cloud-Delivered Firewall Management Center registers with the Smart Software Manager, and then assigns licenses for each managed device. Devices do not register directly with the Smart Software Manager.

A physical Cloud-Delivered Firewall Management Center does not require a license for its own use.

Periodic communications with the Smart Software Manager

You use a Product Instance Registration Token to register the Cloud-Delivered Firewall Management Center with the Smart Software Manager. The Smart Software Manager issues an ID certificate for communication between the Cloud-Delivered Firewall Management Center and the Smart Software Manager. This certificate is valid for one year, although it will be renewed every six months. If an ID certificate expires (after a year with no communication), the Cloud-Delivered Firewall Management Center may be removed from your account.

The Cloud-Delivered Firewall Management Center communicates with the Smart Software Manager on a periodic basis. If you make changes in the Smart Software Manager, you can refresh the authorization on the Cloud-Delivered Firewall Management Center so the changes immediately take effect. You also can wait for the Cloud-Delivered Firewall Management Center to communicate as scheduled. Your Cloud-Delivered Firewall Management Center must either have direct internet access to the Smart Software Manager.

In non-airgapped deployments, normal license communication occurs every 30 days. With the grace period, your Cloud-Delivered Firewall Management Center will operate for up to 90 days without contacting the Smart Software Manager. Ensure that the Cloud-Delivered Firewall Management Center contacts the Smart Software Manager within 90 days, or the Cloud-Delivered Firewall Management Center will revert to an unregistered state.

If your device is unable to communicate due to network issues, it continues working for up to 90 days. Once communication resumes and the ID certificate is validated, license entitlement is restored.

Cloud-Delivered Firewall Management Center licensing and registration

Register Cloud-Delivered Firewall Management Center with Cisco Smart Licensing the same way you would for On-Premises Firewall Management Center.

Overview

Cloud-Delivered Firewall Management Center is included with the base subscription and trial for Security Cloud Control Firewall Management. You do not need to purchase a separate license for Cloud-Delivered Firewall Management Center. The base subscription or trial covers Cloud-Delivered Firewall Management Center, enabling you to manage Firewall Threat Defense devices through the cloud.

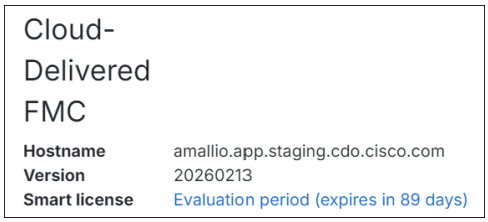
Cloud-Delivered Firewall Management Center is provisioned with a 90-day evaluation license upon deployment.

Important We strongly recommend registering Cloud-Delivered Firewall Management Center with Cisco Smart Software Manager immediately after deployment to access all features, including those requiring export encryption such as RAVPN and SSL VPN. Early registration prevents feature limitations and deployment interruptions.

Evaluation license

- Features requiring export encryption, such as Remote Access VPN, SSL VPN, and other encryption-dependent options, remain disabled until you register Cloud-Delivered Firewall Management Center with Cisco Smart Software Manager (CSSM).
- After the 90-day evaluation period ends, you can continue onboarding Firewall Threat Defense devices, but manually triggered or scheduled deployments are blocked until registration with CSSM is completed.
- Security Cloud Control Firewall Management sends alert notifications as the evaluation license nears expiration.

Note The **Smart license** field displays a warning when the evaluation license is expiration.



When the evaluation period ends, Cloud-Delivered Firewall Management Center is unregistered. Registering Cloud-Delivered Firewall Management Center with Cisco Smart Software Manager converts from evaluation mode to a fully licensed state.

Click the evaluation warning link to open the smart licensing registration page and register Cloud-Delivered Firewall Management Center with Cisco Smart Software Manager.

Registration with Cisco Smart Software Manager (CSSM)

- Connect Cloud-Delivered Firewall Management Center to your smart licensing account, following the same process as On-Premises Firewall Management Center. The high-level procedure is provided in this section.

- No separate license purchase is needed for Cloud-Delivered Firewall Management Center; the license is automatically applied upon registration.
- Registering with smart licensing enables encryption export features.
- Registration ensures compliance and prevents deployment issues.

To learn how to get a Cloud-Delivered Firewall Management Center provisioned on your Security Cloud Control tenant, see Request a Cloud-Delivered Firewall Management Center for your Security Cloud Control Tenant.

Licensing for Firewall Threat Defense devices

- No separate license purchase is needed for Cloud-Delivered Firewall Management Center.
- Each Firewall Threat Defense device managed by Cloud-Delivered Firewall Management Center requires an individual license.
- After registering Cloud-Delivered Firewall Management Center with CSSM, an administrator may apply feature licenses to security devices.

Step-by-step registration

A high-level procedure is provided here. For detailed steps, see Register the Management Center with the Smart Software Manager.

1. Create or Access Your Smart Account:

   Ensure you have a Cisco Smart Software Manager account. If you do not have one, create it at Smart Software Manager.

2. Generate a Registration Token or Use a Registration Token:

   In the Smart Software Manager, use an existing valid registration token or generate a registration token for the virtual account to which you want to register the Cloud-Delivered Firewall Management Center.

   a. Navigate to **Inventory** > **New Token**.

   b. Enter a description and set an expiration (Cisco recommends 30 days).

   c. Create and copy the token.

3. Register Cloud-Delivered Firewall Management Center with the Token:

   a. Log in to the Cloud-Delivered Firewall Management Center interface.

   b. Navigate to **System** > **Licenses** > **Smart Licenses**.

   c. Click **Register** and paste the registration token into the **Product Instance Registration Token** field.

   d. Ensure there are no extra spaces or blank lines before or after the token.

   e. Click **Apply Changes**.

4. Verify Registration:

   After registration, verify the license status under **System** > **Licenses** > **Smart Licenses** to confirm successful registration, license application, and that “Export-Controlled Features” are enabled.

Out-of-compliance state

The Smart License Usage Authorization of Cloud-Delivered Firewall Management Center can become out of compliance in these situations:

- License expiration—When a term-based license of a managed device expires.

In an out-of-compliance state, these effects occur:

- All managed device licenses—Operation is not affected.

After you resolve the licensing problem, the Cloud-Delivered Firewall Management Center will show that it is now in compliance after its regularly scheduled authorization with the Smart Software Manager. To force an authorization, click **Re-Authorize** on the **System** (⚙) > **Licenses** > **Smart Licenses** page.

Unregistered state

The Cloud-Delivered Firewall Management Center can become unregistered in these situations:

- Evaluation mode expiration—Evaluation mode expires after 90 days.
- Manual deregistration of the Cloud-Delivered Firewall Management Center
- Lack of communication with the Smart Software Manager—The Cloud-Delivered Firewall Management Center does not communicate with the Smart Software Manager for 1 year. Note: After 90 days, the Cloud-Delivered Firewall Management Center authorization expires, but it can successfully resume communication within one year to automatically re-authorize. After a year, the ID certificate expires, and the Cloud-Delivered Firewall Management Center is removed from your account so you will have to manually re-register the Cloud-Delivered Firewall Management Center.

In an unregistered state, the Cloud-Delivered Firewall Management Center cannot deploy any configuration changes to devices *for features that require licenses*.

End-user license agreement

The Cisco end-user license agreement (EULA) and any applicable supplemental agreement (SEULA) that governs your use of this product are available from http://www.cisco.com/go/softwareterms.

License types and restrictions

This section describes the types of licenses available for Secure Firewall.

Table 15: Smart Licenses

License You Assign	Duration	Granted Capabilities
Essentials	Perpetual or Subscription **Note** Essentials subscription licenses are supported only on Firewall Threat Defense Virtual .	Except for Specific License Reservation and the Secure Firewall 1200/ 3100 /4200 , Essentials perpetual licenses are automatically assigned with all Firewall Threat Defense s. User and application control Switching and routing NAT For details, see Essentials licenses, on page 281 .
IPS	Subscription	Intrusion detection and prevention File control Security Intelligence filtering For details, see IPS licenses, on page 282
Malware defense	Subscription	Malware defense Secure Malware Analytics File storage (IPS license is a prerequisite for a Malware defense license.) For details, see Malware Defense Licenses, on page 282 and *License Requirements for File and Malware Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide* .
Carrier	Subscription for Firepower 4100/9300, Secure Firewall 3100 /4200 , and Firewall Threat Defense Virtual	Diameter, GTP/GPRS, M3UA, and SCTP inspection For details, see Carrier license, on page 283 .
URL Filtering	Subscription	Category and reputation-based URL filtering For details, see URL Filtering licenses, on page 284 . (IPS license is a prerequisite for a URL Filtering license.)
Export-Controlled Features	Perpetual	Features that are subject to national security, foreign policy, and anti-terrorism laws and regulations; see License for export-controlled functionality, on page 286 .

License You Assign	Duration	Granted Capabilities
Remote Access VPN: • Secure Client Premier • Secure Client Advantage • Secure Client VPN Only	Subscription or perpetual	Remote access VPN configuration. Your account must allow export-controlled functionality to configure remote access VPN. You select whether you meet export requirements when you register the device. The Firewall Threat Defense can use any valid Secure Client license. The available features do not differ based on license type. For more information, see Secure Client licenses, on page 285 and *VPN Licensing* in the *Cisco Secure Firewall Management Center Device Configuration Guide* .



Note Subscription licenses are term-based licenses.

Essentials licenses

The Essentials license allows you to configure and manage firewall devices with essential networking and security capabilities, including switching, routing, high availability, clustering, user control, application control, and database updates.

The Essentials license allows you to:

- Configure your devices to perform switching and routing (including DHCP relay and NAT)
- Configure devices as a high availability pair
- Configure clustering
- Implement user and application control by adding user and application conditions to access control rules
- Update the Vulnerability database (VDB) and geolocation database (GeoDB).
- Download intrusion rules such as SRU/LSP. However, you cannot deploy access control policy or rules that have intrusion policy to the device unless IPS license is enabled.

Secure Firewall 1200/3100/4200

You obtain a Essentials license when you purchase the Secure Firewall 1200/3100/4200.

Other Models

Except in deployments using Specific License Reservation, a Essentials license is automatically added to your account when you register a device to the Cloud-Delivered Firewall Management Center. For Specific License Reservation, you need to add the Essentials license to your account.

Malware Defense Licenses

A Malware defense license lets you perform Malware Defense and Secure Malware Analytics. With this feature, you can use devices to detect and block malware in files transmitted over your network.

To support this feature license, you can purchase the Malware defense (AMP) service subscription as a stand-alone subscription or in combination with IPS (TM) or IPS and URL Filtering (TMC) subscriptions. IPS license is a prerequisite for a Malware defense license.



Note Managed devices with Malware defense licenses enabled periodically attempt to connect to the Secure Malware Analytics Cloud even if you have not configured dynamic analysis. Because of this, the device's Interface Traffic dashboard widget shows transmitted traffic; this is expected behavior.

You configure Malware Defense as part of a file policy, which you then associate with one or more access control rules. File policies can detect your users uploading or downloading files of specific types over specific application protocols. Malware Defense allows you to use local malware analysis and file preclassification to inspect a restricted set of those file types for malware. You can also download and submit specific file types to the Secure Malware Analytics Cloud for dynamic and Spero analysis to determine whether they contain malware. For these files, you can view the network file trajectory, which details the path the file has taken through your network. The Malware Defense license also allows you to add specific files to a file list and enable the file list within a file policy, allowing those files to be automatically allowed or blocked on detection.

Note that a Malware defense license is required only if you deploy Malware Defense and Secure Malware Analytics. Without a Malware defense license, the Cloud-Delivered Firewall Management Center can receive Secure Endpoint malware events and indications of compromise (IOC) from the Secure Malware Analytics Cloud.

Refer to important information at *License Requirements for File and Malware Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

When you disable this license:

- The system stops querying the Secure Malware Analytics Cloud, and also stops acknowledging retrospective events sent from the Secure Malware Analytics Cloud.
- You cannot re-deploy existing access control policies if they include Malware Defense configurations.
- For a very brief time after a Malware defense license is disabled, the system can use existing cached file dispositions. After the time window expires, the system assigns a disposition of **Unavailable** to those files.

If the license expires, your entitlement for the above capabilities ceases and the Cloud-Delivered Firewall Management Center moves to the out-of-compliance state.

IPS licenses

A IPS license allows you to perform intrusion detection and prevention, file control, and Security Intelligence filtering.

License capabilities

- *Intrusion detection and prevention* allows you to analyze network traffic for intrusions and exploits and, optionally, drop offending packets.

- *File control* allows you to detect and, optionally, block users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. *Malware defense*, which requires a Malware defense license, allows you to inspect and block a restricted set of those file types based on their dispositions.
- *Security Intelligence filtering* allows you to block —deny traffic to and from—specific IP addresses, URLs, and DNS domain names, before the traffic is subjected to analysis by access control rules. Dynamic feeds allow you to immediately block connections based on the latest intelligence. Optionally, you can use a "monitor-only" setting for Security Intelligence filtering.

License purchase options

You can purchase a IPS license as a stand-alone subscription (T) or in combination with URL Filtering (TC), Malware defense (TM), or both (TMC).

Disabled license behavior

When you disable this license:

- The Cloud-Delivered Firewall Management Center stops acknowledging intrusion and file events from the affected devices. As a consequence, correlation rules that use those events as a trigger criteria stop firing.
- The Cloud-Delivered Firewall Management Center does not contact the internet for either Cisco-provided or third-party Security Intelligence information.
- You cannot re-deploy existing intrusion policies until you re-enable IPS.

Expired license behavior

If the license expires, your entitlement for the above capabilities ceases and the Cloud-Delivered Firewall Management Center moves to the out-of-compliance state.

Carrier license

The Carrier license enables the inspection of specific protocols used in next-generation mobile and telecom networks. This license supports protocol inspection for Diameter, GTP/GPRS, M3UA, and SCTP protocols across supported devices.

Enabled protocols

- Diameter—Diameter is an Authentication, Authorization, and Accounting (AAA) protocol used in next-generation mobile and fixed telecom networks such as EPS (Evolved Packet System) for LTE (Long Term Evolution) and IMS (IP Multimedia Subsystem). It replaces RADIUS and TACACS in these networks.
- GTP/GPRS—GPRS Tunneling Protocol (GTP) is used in GSM, UMTS, and LTE networks for general packet radio service (GPRS) traffic. GTP provides a tunnel control and management protocol to provide GPRS network access for a mobile station by creating, modifying, and deleting tunnels. GTP also uses a tunneling mechanism for carrying user data packets.
- M3UA—MTP3 User Adaptation (M3UA) is a client/server protocol that provides a gateway to the Signaling System 7 (SS7) network for IP-based applications that interface with the SS7 Message Transfer

Part 3 (MTP3) layer. M3UA makes it possible to run the SS7 User Parts (such as ISUP) over an IP network.

- SCTP—Stream Control Transmission Protocol (SCTP) is a transport-layer protocol that supports the SS7 protocol over IP networks. It supports the 4G LTE mobile network architecture. SCTP can handle multiple simultaneous streams, multiplexed streams, and provides more security features.



Note After you enable this license on a device, use a FlexConfig policy to enable the protocol inspection.

The Carrier license PIDs are available per family and not per device model. You can enable this license for each device either in the evaluation mode or with a Smart License.

The Carrier license for Firepower 4100/9300, Secure Firewall 3100/4200, and Firewall Threat Defense Virtual is term-based. This license also supports Specific License Reservation.

Supported devices

- Cisco Secure Firewall 3100 Series
  - Secure Firewall 3110
  - Secure Firewall 3120
  - Secure Firewall 3130
  - Secure Firewall 3140
- Cisco Secure Firewall 4100 Series
  - Firepower 4112
  - Firepower 4115
  - Firepower 4125
  - Firepower 4145
- Cisco Secure Firewall 4200 Series
  - Secure Firewall 4215
  - Secure Firewall 4225
  - Secure Firewall 4245
- Firepower 9300
- Firewall Threat Defense Virtual

URL Filtering licenses

The URL Filtering license allows you to write access control rules that determine the traffic that can traverse your network based on URLs requested by monitored hosts, correlated with information about those URLs. To support this feature license, you can purchase the URL Filtering service subscription as a stand-alone

subscription or in combination with IPS (TC) or Threat and Malware defense (TMC) subscriptions. IPS license is a prerequisite for this license.

Tip Without a URL Filtering license, you can specify individual URLs or groups of URLs to allow or block. This option gives you granular, custom control over web traffic, but does not allow you to use URL category and reputation data to filter network traffic.

Although you can add category and reputation-based URL conditions to access control rules without a URL Filtering license, the Cloud-Delivered Firewall Management Center will not download URL information. You cannot deploy the access control policy until you first add a URL Filtering license to the Cloud-Delivered Firewall Management Center, then enable it on the devices targeted by the policy.

When you disable this license:

- You may lose access to filtering network traffic based on the URL category and reputation ACP rules. The manual URL filtering options will continue to be supported.
- Access control rules with URL conditions immediately stop filtering URLs.
- Your Cloud-Delivered Firewall Management Center can no longer download updates to URL data.
- You cannot re-deploy existing access control policies if they include rules with category and reputation-based URL conditions.

If the license expires, your entitlement for the above capabilities ceases and the Cloud-Delivered Firewall Management Center moves to the out-of-compliance state.

Secure Client licenses

You can configure remote access VPN using the Secure Client and standards-based IPSec/IKEv2.

To enable remote access VPN, you must purchase and enable one of these licenses: Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only. You can select Secure Client Advantage and Secure Client Premier if you have both licenses and you want to use them both. The Secure Client VPN Onlylicense cannot be used with **Apex** or **Plus**. The Secure Client license must be shared with the Smart Account. For more instructions, refer to http://www.cisco.com/c/dam/en/us/products/collateral/security/anyconnect-og.pdf.

You cannot deploy the remote access VPN configuration to the device if the specified device does not have the entitlement for a minimum of one of the specified Secure Client license types. If the registered license moves out of compliance or entitlements expire, the system displays licensing alerts and health events.

While using remote access VPN, your Smart Account must have the export controlled features (strong encryption) enabled. The Firewall Threat Defense requires strong encryption (which is higher than DES) for successfully establishing remote access VPN connections with Secure Clients.

You cannot deploy remote access VPN if these conditions are true:

- Smart Licensing on the Cloud-Delivered Firewall Management Center is running in evaluation mode.
- Your Smart Account is not configured to use export-controlled features (strong encryption).

License for export-controlled functionality

This topic provides information about software features that are subject to export control laws and regulations, and describes how to determine and enable export-controlled functionality for your system.

Features that require export-controlled functionality

Certain software features are subject to national security, foreign policy, and anti-terrorism laws and regulations. These export-controlled features include:

- Security certifications compliance
- Remote access VPN
- Site-to-site VPN with strong encryption
- SSH platform policy with strong encryption
- SSL policy with strong encryption
- Functionality such as SNMPv3 with strong encryption

Verify export-controlled functionality in your system

To determine whether export-controlled functionality is currently enabled for your system: Choose **System** (⚙) > **Licenses** > **Smart Licenses** and see if **Export-Controlled Features** displays **Enabled**.

Enable export-controlled functionality

If **Export-Controlled Features** shows **Disabled** and you want to use features that require strong encryption, there are two ways to enable strong cryptographic features. Your organization may be eligible for one or the other (or neither), but not both.

- If there is *no* option to enable export-controlled functionality when you generate a new Product Instance Registration Token in the Smart Software Manager, contact your account representative.
- If the option "Allow export-controlled functionality on the products registered with this token" appears when you generate a new Product Instance Registration Token in the Smart Software Manager, make sure you check it before generating the token.

  If you did not enable export-controlled functionality for the Product Instance Registration Token that you used to register the Cloud-Delivered Firewall Management Center, then you must deregister and then re-register the Cloud-Delivered Firewall Management Center using a new Product Instance Registration Token with export-controlled functionality enabled.

If you registered devices to the Cloud-Delivered Firewall Management Center in evaluation mode or before you enabled strong encryption on the Cloud-Delivered Firewall Management Center, reboot each managed device to make strong encryption available. In a high availability deployment, the active and standby devices must be rebooted together to avoid an Active-Active condition.

The entitlement is perpetual and does not require a subscription.

More information

For general information about export controls, see https://www.cisco.com/c/en/us/about/legal/global-export-trade.html.

Firewall Threat Defense Virtual Licenses

Learn about the performance-tiered license entitlements for Firewall Threat Defense Virtual, including supported vCPU and memory configurations, compatible VM resource footprints, and RA VPN session limits by entitlement tier.

Any Firewall Threat Defense Virtual license can be used on any supported Firewall Threat Defense Virtual vCPU/memory configuration. This allows Firewall Threat Defense Virtual customers to run on a wide variety of VM resource footprints. This also increases the number of supported AWS and Azure instances types. When configuring the Firewall Threat Defense Virtual VM, the maximum supported number of cores (vCPUs) is 16 ; and the maximum supported memory is 32 GB RAM .

Performance Tiers for Firewall Threat Defense Virtual Smart Licensing

Session limits for RA VPNs are determined by the installed Firewall Threat Defense Virtual platform entitlement tier, and enforced via a rate limiter. The following table summarizes the session limits based on the entitlement tier and rate limiter.

Table 16: Firewall Threat Defense Virtual Licensed Feature Limits Based on Entitlement

Performance Tier	Device Specifications (Core/RAM)	Rate Limit	RA VPN Session Limit
FTDv5, 100Mbps	4 core/8 GB	100Mbps	50
FTDv10, 1Gbps	4 core/8 GB	1Gbps	250
FTDv20, 3Gbps	4 core/8 GB	3Gbps	250
FTDv30, 5Gbps	8 core/16 GB	5Gbps	250
FTDv50, 10Gbps	12 core/24 GB	10Gbps	750
FTDv100, 16Gbps	16 core/32 GB	16Gbps	10,000
FTDvU	32 core/64 GB	Not limited	20,000
FTDvU	64 core/128 GB	Not limited	32,000

Firewall Threat Defense Virtual Performance Tier Licensing Guidelines and Limitations

Review and follow these guidelines and limitations when licensing your Firewall Threat Defense Virtual device.

- The Firewall Threat Defense Virtual supports performance-tiered licensing that provides different throughput levels and VPN connection limits based on deployment requirements.
- Any Firewall Threat Defense Virtual license can be used on any supported Firewall Threat Defense Virtual core/memory configuration. This allows the Firewall Threat Defense Virtual customers to run on a wide variety of VM resource footprints.
- You can select a performance tier when you deploy the Firewall Threat Defense Virtual , whether your device is in evaluation mode or is already registered with Smart Software Manager

Note Make sure your Smart Licensing account contains the available licenses you need. It's important to choose the tier that matches the license you have in your account. If you are upgrading your Firewall Threat Defense Virtual to Version 7.0 and later, you can choose **FTDv - Variable** to maintain your current license compliance. Your Firewall Threat Defense Virtual continues to perform with session limits based on your device capabilities (number of cores/RAM).

- The default performance tier is FTDv50 when deploying a new Firewall Threat Defense Virtual device, or when provisioning the Firewall Threat Defense Virtual using the REST API.
- Essentials licenses are subscription-based and mapped to performance tiers. Your virtual account needs to have the Essentials license entitlements for the Firewall Threat Defense Virtual devices, as well as for IPS , Malware Defense , and URL Filtering licenses.
- Each HA peer consumes one entitlement, and the entitlements on each HA peer must match, including Essentials license.
- A change in performance tier for an HA pair should be applied to the primary peer.
- You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.
- Universal PLR licensing is applied to each device in an HA pair separately. The secondary device will not automatically mirror the performance tier of the primary device. It must be updated manually.

Edit performance tier of Firewall Threat Defense Virtual

The Firewall Threat Defense Virtual supports performance-tiered licensing that provides different throughput levels and VPN connection limits based on deployment requirements. You can edit the performance tier using this procedure.

Before you begin

Ensure that you review Firewall Threat Defense Virtual Performance Tier Licensing Guidelines and Limitations , on page 287.

Procedure

Step 1 Choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Step 2 Click **Edit Performance Tier**.

Step 3 Choose a tier from the **Target Performance Tier** drop-down list.

Step 4 From the **Available Devices** list, select the required devices and click **Add**.

Step 5 Click **Apply**.

Prerequisites for licensing

This topic lists the prerequisites, supported domains, and user roles required for licensing.

General prerequisites

- Make sure NTP is configured on the Cloud-Delivered Firewall Management Center and managed devices. Time must be synchronized for registration to succeed.
- For a Firepower 4100/9300, you must configure NTP on the chassis using the same NTP server for the chassis as for the Cloud-Delivered Firewall Management Center.

Supported domains

Global, except where indicated.

User roles

Admin

Prerequisites for licensing in high availability, clustering, and multi-instance

This section describes the licensing requirements for device High Availability.

FTD Services does not support clustering or multi-instance deployments.

License requirements for Firewall Threat Defense device high availability

Both Firewall Threat Defense units in a high availability configuration must have the same licenses.

High availability configurations require two license entitlements: one for each device in the pair.

Before high availability is established, it does not matter which licenses are assigned to the secondary/standby device. During high availability configuration, the Cloud-Delivered Firewall Management Center releases any unnecessary licenses assigned to the standby unit and replaces them with identical licenses assigned to the primary/active unit.

For example, if the active unit has a Essentials license and a IPS license, and the standby unit has only a Essentials license, the Cloud-Delivered Firewall Management Center communicates with the Smart Software Manager to obtain an available IPS license from your account for the standby unit. If your license account does not include enough purchased entitlements, your account becomes Out-of-Compliance until you purchase the correct number of licenses.

License for device clusters

Device clusters require specific licensing considerations for proper operation and feature functionality.

Performance tier license requirements

Each Firewall Threat Defense Virtual cluster node requires the same performance tier license. We recommend using the same number of CPUs and memory for all members, or else performance will be limited on all nodes

to match the least capable member. The throughput level will be replicated from the control node to each data node so they match.

Feature license assignment

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the Cloud-Delivered Firewall Management Center, you can specify the feature licenses you want to use for the cluster. Before you create the cluster, it doesn't matter which licenses are assigned to the data nodes; the license settings for the control node are replicated to each of the data nodes. You can modify licenses for the cluster by clicking **Edit Licenses** in **System** (⚙) > **Licenses** > **Smart Licenses** or choosing **Devices** > **Device Management**, clicking **Edit** (✎) for the cluster, and then in the **License** area, clicking **Edit** (✎).

Note If you add the cluster before the Cloud-Delivered Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Cloud-Delivered Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Note APEX, PLUS, and VPN-Only licenses are not supported for clustering. Before you create a cluster, remove unsupported license assignments from the devices that will become cluster nodes.

Create a Cisco account

You need a Cisco account to access Smart Account functionality and license any Cisco products.

You must have a Cisco account to request a Smart Account and license any Cisco products.

Procedure

Step 1 Open the URL https://id.cisco.com/signin/register to create a new account.

Step 2 Enter all the required fields to create an account.

Step 3 Click **Register**.

An email with an activation code is sent to verify your email address.

Note
If you haven't received an email yet, send an email to the registration support team at web-help@cisco.com.

Step 4 In the **Verify with your email** page, enter the activation code to complete the registration process and click **Verify**.

After successful registration, you will be redirected to the login page.

What to do next

Enter the newly created account details on the login page to request a Smart Account. Refer to Create a smart account and add licenses, on page 291.

Create a smart account and add licenses

Create a Smart Account to centrally manage and track your Cisco software licenses across your organization.

You should set up this account before you purchase licenses. Your account representative or reseller may have set up a Smart Account on your behalf. If so, obtain the necessary information to access the account from that person instead of using this procedure, then verify that you can access the account.

Before you begin

You must create a new Cisco account if you don't already have one. For instructions, refer to Create a Cisco Account.

For general information about Smart Accounts, refer to http://www.cisco.com/go/smartaccounts.

Procedure

Step 1 Go to the Create a Smart Account page.

You are prompted to log in with your Cisco account. In the **Create a Smart Account** page, your basic account information is displayed.

Step 2 Click the **My Account** icon appearing in the top right corner and click **Manage Profile**.



Step 3 Click **Personal** and configure these parameters:

a) In the **Your Company Details** section, click **Edit**.

b) In the **Company or organization** field, type your organization name.

c) If your company information is already present in our database, it will appear in the list. You can select your company.

In the **Address** drop-down list, select the address of your company.

d) If your company is not listed in our database, you can continue to enter your company information in the **Company or organization** field.

1. In the **Address** drop-down list, click the drop-down arrow and click **Add New Address**.
2. You can select one of the following **Address Type** options:
   - **Company/Organization**: To provide your organization's address. Cisco verifies this address. If the address and company name cannot be validated with the country, you may not be able to proceed. Therefore, you must ensure that the correct address is provided.
   - **Personal**: To provide your personal address.

e) Enter all the mandatory fields associated with your company and click **Update**.

The **Your Company Details** section displays the company details you entered.

A success message appears if your company details are validated.

Step 4 Click **Update**.

A success message appears if your company details are validated.

Step 5 Open the **Create a Smart Account** page, which was opened in the previous tab. Refresh the page if changes are not reflecting.

Alternatively, you can open this page using the https://software.cisco.com/software/company/smartaccounts/home?route=module/accountcreation URL and login using your credentials.

Step 6 Click **Create Account**.

The **Account Summary** page displays your account details.

Step 7 Click **Done**.

Step 8 Wait for an email telling you that your Smart Account is ready to set up. When it arrives, click the link it contains, as directed.

Step 9 Make sure your Smart Licensing account contains the available licenses you need.

What to do next

To configure Smart License using the Smart Software Manager, refer to Configure Smart Licensing, on page 292.

Configure Smart Licensing

This section describes how to use Smart Licensing using the Smart Software Manager or the Smart Software Manager On-Prem.

Register the Cloud-Delivered Firewall Management Center for Smart Licensing

You can register the Cloud-Delivered Firewall Management Center directly to the Smart Software Manager over the internet, or when using an air-gapped network, with the Smart Software Manager On-Prem.

Register the Cloud-Delivered Firewall Management Center with the Smart Software Manager

Register the Cloud-Delivered Firewall Management Center with the Smart Software Manager to enable Smart Licensing functionality and proper license management.

Before you begin

- Make sure your Smart Licensing account contains the available licenses you need.

  When you bought your device from Cisco or a reseller, your licenses should have been linked to your Smart Account. However, if you need to add licenses yourself, see Cisco Commerce Workspace.

- Ensure that the Cloud-Delivered Firewall Management Center can reach the Smart Software Manager at smartreceiver.cisco.com.
- Configure NTP. During registration, a key exchange occurs between the Smart Agent and the Smart Software Manager, so time must be in sync for proper registration.

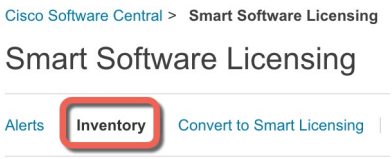
  For the Firepower 4100/9300, configure NTP on the chassis using the same NTP server as the Cloud-Delivered Firewall Management Center.

- If your organization has multiple Cloud-Delivered Firewall Management Centers, make sure each one has a unique name that clearly identifies and distinguishes it from others that may be registered to the same virtual account. This name is critical for managing your Smart License entitlements and ambiguous names will lead to problems later.

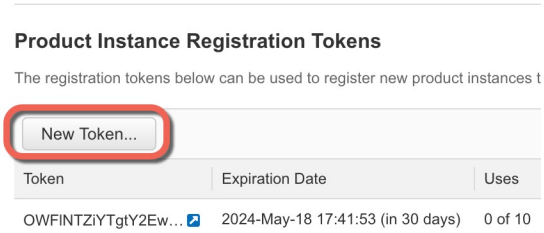
Procedure

Step 1 In the Smart Software Manager, request and copy a registration token for the virtual account to which you want to add this device.

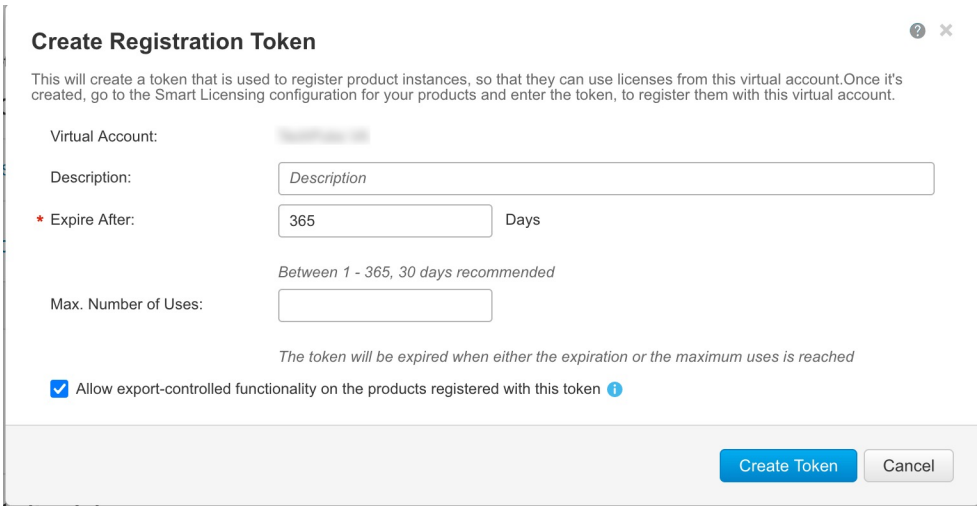
a) Click **Inventory**.



b) On the **General** tab, click **New Token**.



c) In the **Create Registration Token** dialog box enter these settings, and then click **Create Token**:



- **Description**
- **Expire After**—Cisco recommends 30 days.
- **Max. Number of Uses**
- **Allow export-controlled functionality on the products registered with this token**—Enables the export-compliance flag if you are in a country that allows for strong encryption. You must select this option now if you plan to use this functionality. If you enable this functionality later, you will need to re-register your device with a new product key and reload the device. If you do not see this option, your account does not support export-controlled functionality.

The token is added to your inventory.

d) Click the arrow icon to the right of the token to open the **Token** dialog box so you can copy the token ID to your clipboard. Keep this token ready for later in the procedure when you need to register the Firewall Threat Defense.

Figure 32: View token

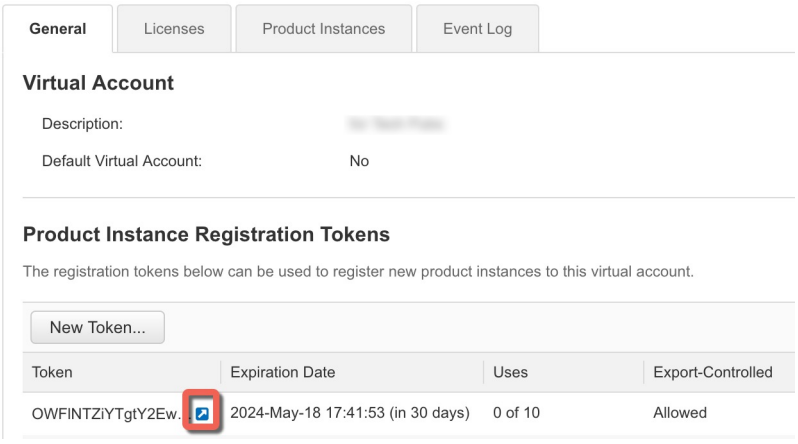
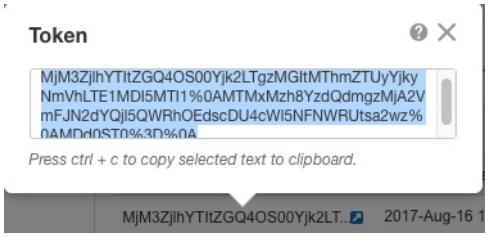


Figure 33: Copy token



Step 2 In the Cloud-Delivered Firewall Management Center, choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Step 3 Click **Register**.

Step 4 Paste the token you generated from Smart Software Manager into the **Product Instance Registration Token** field.

Make sure there are no empty spaces or blank lines at the beginning or end of the text.

Step 5 Click **Apply Changes**.

What to do next

- Add a device to the Cloud-Delivered Firewall Management Center; see *Add a Device to the Cloud-Delivered Firewall Management Center* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Assign licenses to devices

You can assign most licenses when you register a device to the Cloud-Delivered Firewall Management Center. You can also assign licenses per device, or for multiple devices.

Assign licenses to a single device

This task allows you to enable or disable specific licensed features on a managed device by assigning or removing licenses as needed.

You can enable or disable specific licensed features on a managed device by assigning or removing licenses as needed. Although there are some exceptions, you cannot use the features associated with a license if you disable it on a managed device.

Note For container instances on the same security module/engine, you apply the license to each instance; note that the security module/engine consumes only one license per feature for all instances on the security module/engine.

Note For the Firewall Threat Defense cluster, you apply the licenses to the cluster as a whole; note that each unit in the cluster consumes a separate license per feature.

Before you begin

You must have Admin or Network Admin privileges to perform this task. When operating with multiple domains, you must do this task in leaf domains.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to assign or disable a license, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 Next to the **License** section, click **Edit** ( ).

Step 5 Check or clear the appropriate check boxes to assign or disable licenses for the device.

Step 6 Click **Save**.

Step 7 Deploy configuration changes.

What to do next

Verify license status:

1. Choose **System** ( ) > **Licenses** > **Smart Licenses**.
2. Enter the hostname or IP address of the device into the filter at the top of the Smart Licenses table
3. Verify that only a green circle with a **Check Mark** ( ) appears for each device, for each license type.

   If you see any other icon, hover over the icon for more information.

Assign licenses to multiple managed devices

Devices managed by the Cloud-Delivered Firewall Management Center obtain their licenses via the Cloud-Delivered Firewall Management Center, not directly from the Smart Software Manager.

Note For container instances on the same security module/engine, you apply the license to each instance; note that the security module/engine consumes only one license per feature for all instances on the security module/engine.

Note For the Firewall Threat Defense cluster, you apply the licenses to the cluster as a whole; note that each unit in the cluster consumes a separate license per feature.

Use this procedure to enable licensing on multiple devices at once.

Procedure

Step 1 Choose **System** (⚙) > **Licenses** > **Smart Licenses** or **Specific Licenses**.

Step 2 Click **Edit Licenses**.

Step 3 For each type of license you want to add to a device:

a) Click the tab for that type of license.
b) Click a device in the list on the left.
c) Click **Add** to move that device to the list on the right.
d) Repeat for each device to receive that type of license.

For now, don't worry about whether you have licenses for all of the devices you want to add.

e) Repeat this subprocedure for each type of license you want to add.
f) To remove a license, click the **Delete** ( ) next to the device.
g) Click **Apply**.

You can select a cluster and assign any license to all nodes of a cluster.

What to do next

Verify that your licenses are correctly installed. Follow the procedure in Monitor Smart Licensing, on page 298.

Manage Smart Licensing

This section describes how to manage Smart Licensing.

Deregister the Cloud-Delivered Firewall Management Center

Deregister your Cloud-Delivered Firewall Management Center from the Smart Software Manager to release all of the license entitlements back to your Smart Account so they can be used for other devices.

Deregister if you need to decommission the Cloud-Delivered Firewall Management Center or reimage it.

Refer to Unregistered state, on page 279 for more information about license enforcement in an unregistered state.

Procedure

Step 1 Choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Step 2 Click **Deregister** ( ).

Monitor Smart License status

The **Smart License Status** section of the **System** (⚙) > **Licenses** > **Smart Licenses** page provides an overview of license usage on the Cloud-Delivered Firewall Management Center.

Usage authorization

Possible status values are:

- **In-compliance** (🟢) — All licenses assigned to managed devices are in compliance and the Cloud-Delivered Firewall Management Center is communicating successfully with the Smart Software Manager.
- **License is in compliance but communication with licensing authority has failed**— Device licenses are in compliance, but the Cloud-Delivered Firewall Management Center is not able to communicate with the Cisco licensing authority.
- **Out-of-compliance icon or unable to communicate with License Authority**— One or more managed devices is using a license that is out of compliance, or the Cloud-Delivered Firewall Management Center has not communicated with the Smart Software Manager in more than 90 days.

Product registration

Specifies the last date when the Cloud-Delivered Firewall Management Center contacted the Smart Software Manager and registered.

Assigned virtual account

Specifies the Virtual Account under the Smart Account that you used to generate the Product Instance Registration Token and register the Cloud-Delivered Firewall Management Center. If this deployment is not associated with a particular virtual account within your Smart Account, this information is not displayed.

Export-Controlled features

If this option is enabled, you can deploy restricted features. For details, see License for export-controlled functionality, on page 286.

Cisco success network

Specifies whether you have enabled Cisco Success Network for the Cloud-Delivered Firewall Management Center. If this option is enabled, you provide usage information and statistics to Cisco which are essential to provide you with technical support. This information also allows Cisco to improve the product and make you aware of unused available features so that you can maximize the value of the product in your network.

Monitor Smart Licensing

To view the license status for the Cloud-Delivered Firewall Management Center and its managed devices, use the Smart Licenses page.

For each type of license in your deployment, the page lists the total number of licenses consumed, whether the license is in compliance or out of compliance, the device type, and the domain and group where the device is deployed. You can also view the Cloud-Delivered Firewall Management Center's Smart License Status. Container instances on the same security module/engine only consume one license per security module/engine.

Therefore, even though the Cloud-Delivered Firewall Management Center lists each container instance separately under each license type, the number of licenses consumed for feature license types will only be one.

Other than the **Smart Licenses** page, there are a few other ways you can view licenses:

- The **Product Licensing** dashboard widget provides an at-a-glance overview of your licenses.
- The **Device Management** page (**Devices** > **Device Management**) lists the licenses applied to each of your managed devices.
- The **Smart License Monitor** health module communicates license status when used in a health policy.

Procedure

Step 1 Choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Step 2 In the **Smart Licenses** table, click the arrow at the left side of each **License Type** folder to expand that folder.

Step 3 In each folder, verify that each device has a green circle with a **Check Mark** (✓) in the **License Status** column.

If all devices show a green circle with a **Check Mark** (✓), your devices are properly licensed and ready to use.

If you see any License Status other than a green circle with a **Check Mark** (✓), hover over the status icon to view the message.

What to do next

- If you had any devices that did not have a green circle with a **Check Mark** (✓), you may need to purchase more licenses.

Troubleshoot Smart Licensing

This topic lists some common Smart Licensing issues and provides troubleshooting steps.

Expected licenses do not appear in my smart account

If the licenses you expect to see are not in your Smart Account, try the following:

- Make sure they are not in a different Virtual Account. Your organization's license administrator may need to assist you with this.
- Check with the person who sold you the licenses to be sure that transfer to your account is complete.

Unexpected out-of-compliance notification or other error

- If a device is already registered to a different Cloud-Delivered Firewall Management Center, you need to deregister the original Cloud-Delivered Firewall Management Center before you can license the device

under a new Cloud-Delivered Firewall Management Center. See Deregister the Cloud-Delivered Firewall Management Center, on page 297.

- Check if the term of the subscription license has expired.

Troubleshoot other issues

For solutions to other common issues, see Use FMC and FTD Smart License Registration and Common Issues to Troubleshoot.

Configure Legacy Cloud-Delivered Firewall Management Center PAK-Based Licenses

The Cloud-Delivered Firewall Management Center supports either a Smart license or a legacy PAK (Product Activation Key) license for its platform license. This procedure describes how to apply a PAK-based license.

After re-registration of your Smart Account, you must manually add the classic licenses for all classic devices.

Before you begin

- Make sure you have the product activation key (PAK) from the Software Claim Certificate that Cisco provided when you purchased the license. If you have a legacy, pre-Cisco license, contact Support.

Procedure

Step 1 Note the license key for the Cloud-Delivered Firewall Management Center.

The license key uniquely identifies the Cloud-Delivered Firewall Management Center in the Smart Software Manager. It is composed of a product code (for example, 66) and the MAC address of the management port (eth0) of the Cloud-Delivered Firewall Management Center; for example, 66:00:00:77:FF:CC:88.

a) Choose **System** (⚙) > **Licenses** > **Classic Licenses**.
b) Click **Add New license**.
c) Note the value in the **license Key** field at the top of the **Add Feature license** dialog.

Step 2 Choose **System** (⚙) > **Licenses** > **Classic Licenses**.

Step 3 Click **Add New license**.

Step 4 Continue as appropriate:

- If you have already obtained the license text, skip to Step 8.
- If you still need to obtain the license text, go to the next step.

Step 5 Click **Get license** to open the License Registration Portal.

Note
If you cannot access the Internet using your current computer, switch to a computer that can, and browse to http://cisco.com/go/license.

Step 6 Generate a license from the PAK in the License Registration Portal: https://cisco.com/go/license.

This step requires the PAK you received during the purchase process, as well as the LICENSE key for the Cloud-Delivered Firewall Management Center.

For more information on using this portal, see:

https://slexui.cloudapps.cisco.com/SWIFT/LicensingUI/Quickstart

You will need your account credentials in order to access these links.

Step 7 Copy the license text from either the License Registration Portal display, or the email the License Registration Portal sends you.

Important
The licensing text block in the portal or email message may include more than one license. Each license is bounded by a BEGIN LICENSE line and an END LICENSE line. Make sure that you copy and paste only one license at a time.

Step 8 Return to the **Add Feature License** page in the management center virtual's web interface.

Step 9 Paste the license text into the **License** field, click **Verify License**, and click **Submit License**.

If the license is invalid, make sure that you correctly copied the license text.

Additional information about licensing

For additional information to help resolve common licensing questions, see these documents:

- Licensing FAQ
- License Roadmap

Overview

Some secure deployment environments may restrict or disable access to the Internet. In such environments, the Firepower Management Center is unable to communicate with the License Authority either through the Cisco Smart Software Manager (CSSM) or Smart Software Satellite Server. These deployments can use Universal Permanent License Reservation (UPLR). Permanent licenses do not require access to the internet or the Cisco licensing authority.

The Universal Permanent License permits perpetual, unlimited use of supported Firepower products, if your account on Cisco Smart Software Manager is authorized.

Features that require access to the internet, such as URL Lookups or contextual cross-launch to public web sites, will not work.

Cisco does not collect web analytics or telemetry data for deployments that use Universal Permanent Licensing.

Note Various names are used at Cisco for this functionality, including: Universal Permanent License Reservation, Universal License Reservation, Permanent License Reservation, UPLR, ULR, and PLR. These terms may also be used at Cisco to refer to similar but not necessarily identical licensing models. In this document, PLR and UPLR both mean Universal Permanent License Reservation.

Requirements

Usage of Universal Permanent License Reservation requires approval and authorization from Cisco.

Work with your account representative to determine whether you qualify for this licensing option.

Cisco TAC cannot assist with obtaining approval for this licensing; please do not contact TAC about obtaining approval.

Supported Platforms and Versions

Product	Version	Platform/Hosting Environment
Firepower Management Center	6.3 or later	All
Firepower Management Center Virtual	6.3 or later	All
Firepower Threat Defense	6.1.0 or later	All
Classic devices: ASA with FirePOWER Services NGIPSv 7000- and 8000-Series	Not supported.	

How to Implement Universal Permanent Licensing

	Do This	More Information
Step 1	Meet the prerequisites for this feature.	Prerequisites, on page 303
Step 2	Verify that your Smart Account is ready to deploy Universal Permanent License Reservation.	Verify that your Smart Account is Ready to Deploy Universal Permanent License Reservation, on page 304
Step 3	Enable Universal Permanent License Reservation using the Firepower Management Center shell	Enable the Permanent Licensing Menu Option, on page 305

	Do This	More Information
Step 4	Generate a Reservation Request Code from the Firepower Management Center	Generate a Reservation Request Code from the Firepower Management Center, on page 306
Step 5	Use the Reservation Request Code to generate a Reservation Authorization Code from Cisco Smart Software Manager	Generate a Reservation Authorization Code from Cisco Smart Software Manager, on page 306
Step 6	Enter the Reservation Authorization Code into the Firepower Management Center	Enter the Reservation Authorization Code into the Firepower Management Center, on page 307
Step 7	Apply Permanent Licenses to managed Firepower Threat Devices	Apply Permanent Licenses to Managed Firepower Threat Devices, on page 307
Step 8	(Outside of your Firepower Management Center) Schedule reminders for ongoing maintenance tasks	See the "Maintain Your Air-Gapped Deployment" topic in the Firepower Management Center online help.

Prerequisites

Before you can begin implementing a Universal Permanent License Reservation, you must complete the following:

- If you have not already done so, create your Smart Account.

  For more information, see Create a smart account and add licenses.

- Work with your account representative to obtain approval for Universal Permanent License Reservation for your Firepower products.

- Purchase standard licenses and service subscriptions for your products, sufficient to cover all devices and features you will deploy.

  For descriptions of standard Firepower licenses and service subscriptions, see License types and restrictions.

- Arrange for export-controlled strong cryptographic functionality, if required and if your organization is eligible, and confirm that your account is enabled to use it. Your account representative can assist you with this.

- Confirm with your account representative that the Universal Permanent License is ready for use and reflected in your Smart Account.

- Make sure NTP is configured on the Firepower Management Center and managed devices. Time must be synchronized for registration to succeed. For more information, see the online help in FMC.

  If you are deploying FTD on a Firepower 4100/9300 chassis, you must configure NTP on the Firepower chassis using the same NTP server for the chassis as for the Firepower Management Center.

- If you are currently using standard Smart Licensing on your Secure Firewall Management Center, de-register the Secure Firewall Management Center to release the licenses back to your account. For information, see Deregister the Cloud-Delivered Firewall Management Center.

- (Recommended) Add managed devices to your Secure Firewall Management Center. For instructions, see Register with a new management center. (You can add managed devices at any time, but adding them

now simplifes this process.) You will need to enable the evaluation license in order to do this (under **System > Licenses > Smart Licenses**). Evaluation licensing does not require a connection to the License Authority.

- (Recommended) If you will deploy a Secure Firewall Management Center pair in a high availability configuration, configure that before you assign licenses. (FMCs in a high availability configuration require the same number of licenses as a single FMC.) If you have already deployed licenses to the secondary appliance, de-register licensing from that appliance.

Verify that your Smart Account is Ready to Deploy Universal Permanent License Reservation

In order to prevent problems with deploying your permanent license, complete this procedure before you make any changes in your Cisco Smart Software Manager.

Before you begin

- Complete all Prerequisites, on page 303.
- Make sure you have your Secure Firewall Management Center credentials.

Procedure

Step 1 Sign in to the Secure Firewall Management Center:

https://software.cisco.com/#SmartLicensing-Inventory

Step 2 If applicable, select the correct account from the top right corner of the page.

Step 3 Click **Manage licenses** under the **Smart Software Manager** tile.

Step 4 Click the **Inventory** tab.

Step 5 Click the **Licenses** tab.

Step 6 Verify the following:

- There is a **License Reservation** button.
- In the **License** column, there is a **Cisco FMC Universal License Feature** license.
- In the **Purchased** column for this Cisco FMC Universal License Feature, the number of available licenses matches the number of Secure Firewall Management Center instances (hardware or virtual) that you will deploy.
- There are enough platform and feature licenses for the devices and features you will deploy, including Base licenses and Secure Firewall Management Center Virtual entitlements for your devices.

Step 7 If any of these items is missing or incorrect, contact your account representative to resolve the problem.

Important

Do **not** continue with this process until all the problems are corrected!

Enable the Permanent Licensing Menu Option

Universal License	Supported Devices	Supported Domains	Access
Any	Secure Firewall Management Center	Global only	Admin



Procedure

Step 1 Access the Firepower Management Center console using a USB keyboard and VGA monitor, or use SSH to access the management interface.

Step 2 Log into the Firepower Management Center **admin** account.

FMC versions 6.3 and 6.4: By default, this gives you access to the Linux shell. If the Firepower Management Center CLI is enabled, this gives you access to the command line interface.

FMC versions 6.5 and later: By default, this gives you access to the command line interface.

Step 3 If you are in the command line interface, enter the **expert** command to access the Linux shell.

Step 4 Enable Universal Permanent Licensing by executing the following command:

sudo manage_plr.pl

Example:

```
admin@fmc63betaplr: ~$ sudo manage_plr.pl
Password:
*************** Configuration Utility *************

1 Show PLR Status
2 Enable PLR
3 Disable PLR
0 Exit ************************************************************
Enter choice:
```

Step 5 Enable Permanent License Reservation by selecting option 2.

Step 6 Select option 0 to exit the manage_plr utility.

Step 7 Type **exit** to exit the Linux shell.

Step 8 If applicable, exit the command line interface.

Versions 6.3 and 6.4: If the CLI is enabled, enter **exit**.

Version 6.5 and later: Enter the **exit** command.

Step 9 Verify that you can access the **Permanent License Reservation** page in the Secure Firewall Management Center web interface:

- If the **System** > **Licenses** > **Smart Licenses** page is currently displayed, refresh the page.
- Otherwise, choose **System** > **Licenses** > **Universal Licenses**.

Generate a Reservation Request Code from the Firepower Management Center

Universal License	Supported Devices	Supported Domains	Access
Any	Secure Firewall Management Center	Global only	Admin

Procedure

Step 1 If you are not already viewing the **Universal License Reservation** page in the Firepower Management Center, choose **System** > **Licenses** > **Universal Licenses**.

Step 2 Click **Generate**.

Step 3 Make a note of the **Reservation Request Code**.

Generate a Reservation Authorization Code from Cisco Smart Software Manager

Procedure

Step 1 Go to the Cisco Smart Software Manager:

https://software.cisco.com/#SmartLicensing-Inventory

Step 2 If necessary, select the correct account from the top right of the page.

Step 3 If necessary, click **Inventory**.

(This page may display automatically.)

Step 4 Click the **Licenses** tab.

Step 5 Click **License Reservation**.

Step 6 Enter the code that you generated from Firepower Management Center into the **Reservation Request Code** box.

Step 7 Click **Next**.

Step 8 Select **Cisco FMC Universal License Feature**.

Step 9 Click **Next**.

Step 10 Click **Generate Authorization Code**.

At this point, the license is now in use according to the Smart Software Manager.

Step 11 Copy, download, or note down the Authorization Code in preparation for entering it into the Secure Firewall Management Center.

Enter the Reservation Authorization Code into the Firepower Management Center

Universal License	Supported Devices	Supported Domains	Access
Any	Secure Firewall Management Center	Global only	Admin

Procedure

Step 1 If you are not already viewing the **Permanent License Reservation** page in the Firepower Management Center web interface, choose **System** > **Licenses** > **Universal Licenses**.

Step 2 Enter the generated authorization code in the **Authorization Code** text box.

Step 3 Click **Install**.

Step 4 Verify that the **Permanent License Reservation** page shows the **Usage Authorization** status as authorized.

Apply Permanent Licenses to Managed Firepower Threat Devices

Universal License	Supported Devices	Supported Domains	Access
Any	Secure Firewall Management Center	Global only	Admin

Procedure

Step 1 Choose **System** > **Licenses** > **Universal Licenses**.

Step 2 Click **Edit Licenses**.

Step 3 For each type of license you want to add to a device:

a) Click the tab for that type of license.
b) Click a device in the list on the left.
c) Click **Add** to move that device to the list on the right.
d) Repeat for each device to receive that type of license.

Ensure you have licenses for all of the devices you want to add.

e) Repeat this sub procedure for each type of license you want to add.
f) Click **Apply**.

Step 4 Verify that your licenses are assigned as you expect them, and that each feature for each managed device shows a green circle with a **Check Mark** ( ).

If necessary, see Universal Permanent License Status, on page 308 for information.

What to do next

- If you applied a license with export-controlled functionality enabled, reboot each managed device.

  If devices are configured in a high-availability pair, reboot both devices at the same time to avoid an Active-Active condition.

- Deploy configuration changes to managed devices; see the "Deploy Configuration Changes" topic in the Secure Firewall Management Center online help.

Maintain your air-gapped deployment

Manually obtain and install essential content updates when the Cloud-Delivered Firewall Management Center cannot access the internet.

Refer to:

- Manually update the VDB, on page 265
- Manually update intrusion rules, on page 269
- Manually update the GeoDB, on page 267

Universal Permanent License Status

Universal License Status (on the **System** > **Licenses** > **Universal Licenses** page) provides an overview of license usage on the Firepower Management Center, as described below.

Usage Authorization

Authorized— The product is successfully registered and in compliance after successful installation of the authorization code.

Product Registration

Specifies the date when the Authorization Code was installed for Universal Permanent License.

Export-Controlled Features

Specifies whether export-controlled functionality is enabled for your deployment. If it is not, and you expect it to be, contact Cisco TAC.

Product Instance

The Universally Unique Identifier (UUID) of this Firepower Management Center. This identifies the Firepower Management Center in Cisco Smart Software Manager.

Maintain Your UPLR Deployment

To ensure that your deployment remains effective, you must manually obtain and install updates. For more information, see Maintain your air-gapped deployment.

Deactivate and Return a Universal Permanent License

Universal License	Supported Devices	Supported Domains	Access
Any	Secure Firewall Management Center	Global only	Admin

If you no longer need a permanent license, you must return it to your Smart Account.



Important If you do not follow all of the steps in this procedure, the license remains in an in-use state in Cisco Smart Software Manager and cannot be re-used.

Procedure

Step 1 In the Secure Firewall Management Center web interface, choose **System** > **Licenses** > **Universal Licenses**.

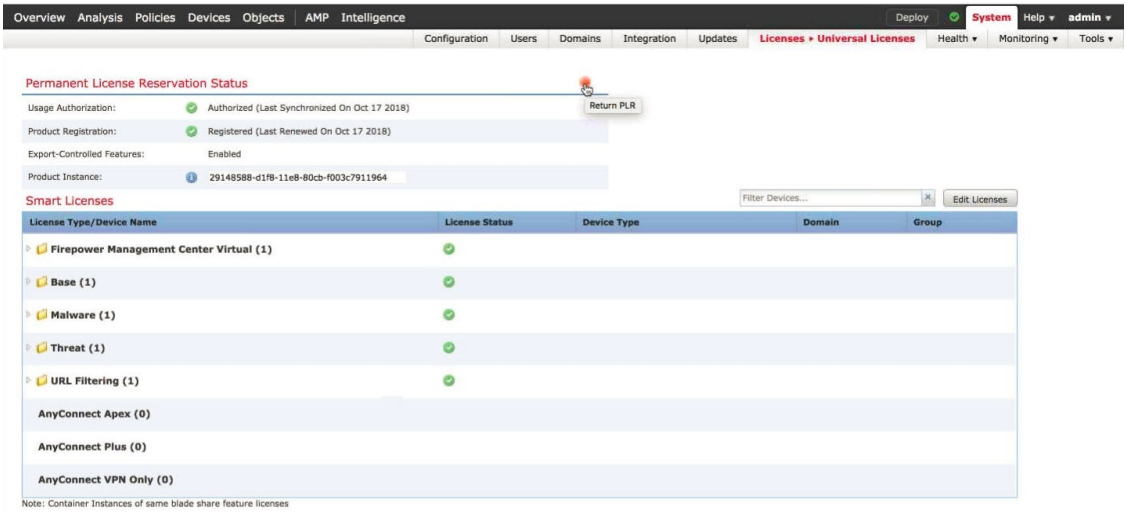
Step 2 Make a note of the **Product Instance** identifier for this Secure Firewall Management Center.

Important
You cannot complete this procedure without this identifier.

Step 3 Generate a return code from Firepower Management Center:

a) Click the **Return PLR** button.

The following figure shows the **Return PLR** button.



Firepower Threat Defense becomes unlicensed and Firepower Management Center moves to the de-registered state.

b) Make a note of the **Return Code**.

Step 4 In Cisco Smart Software Manager, identify the Firepower Management Center appliance to deregister:

a) Go to the Cisco Smart Software Manager:

https://software.cisco.com/#SmartLicensing-Inventory

b) If necessary, click **Inventory**.

(This page may display automatically.)

c) Click **Product Instances**.

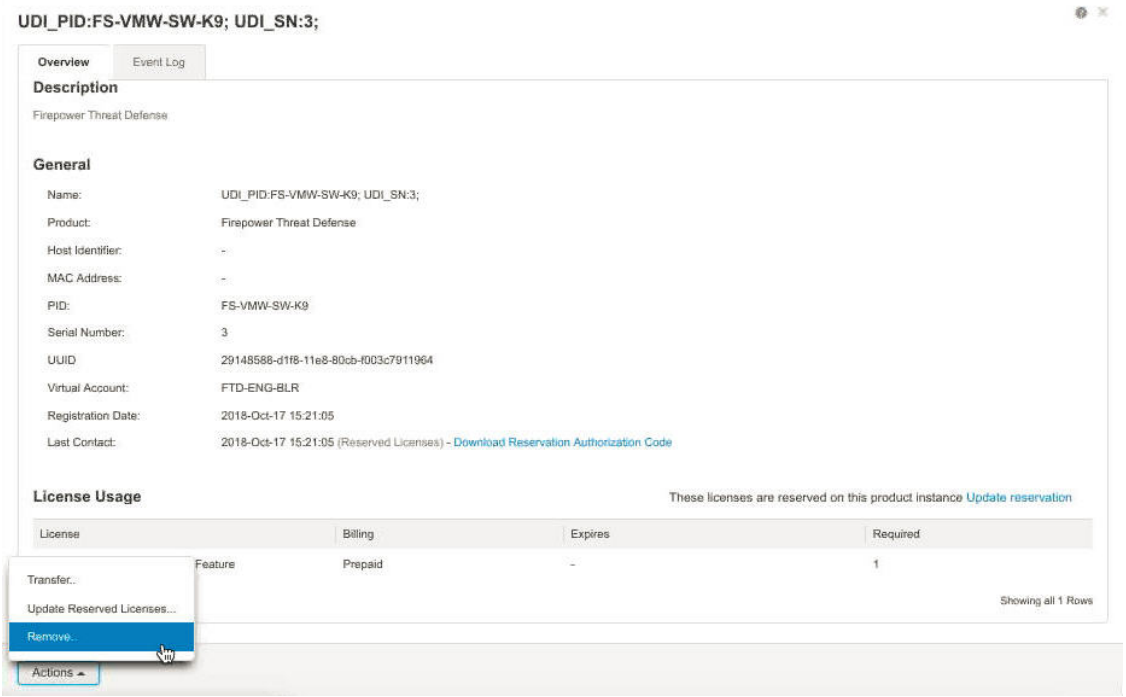
d) Look for a product instance that has **FP** in the **Type** column and a generic SKU (not a hostname) in the **Name** column. You may also be able to use the values in other table columns to help determine which Firepower Management Center is the correct Firepower Management Center. Click the name.

e) Look at the **UUID** and see if it is the UUID of the Firepower Management Center that you are trying to modify.

If not, you must repeat these steps until you find the correct Firepower Management Center.

Step 5 When you have identified the correct Firepower Management Center, return the license to your Smart Account:

a) On the device page, choose **Actions** > **Remove**.



b) Enter the reservation return code that you generated from Secure Firewall Management Center into the **Remove Product Instance** dialog box.

c) Click **Remove Product Instance**.

The Permanent License is returned to the pool of available licenses in the Smart Account.

Step 6 Disable the Universal Permanent License in the Secure Firewall Management Center CLI:

a) Access the Firepower Management Center console using a USB keyboard and VGA monitor, or use SSH to access the management interface.

b) Log in to the Firepower Management Center admin account.

FMC versions 6.3 and 6.4: By default, this gives you access to the Linux shell. If the Firepower Management Center CLI is enabled, this gives you access to the command line interface.

FMC versions 6.5 and later: By default, this gives you access to the command line interface.

c) If you are in the command line interface, enter the **expert** command to access the Linux shell.

d) Execute the following command:

sudo manage_plr.pl

Example:

```
admin@fmc63betaplr: ~$ sudo manage_plr.plnp
Password:
*************** Configuration Utility *************
1 Show PLR Status
2 Enable PLR
3 Disable PLR
0 Exit ************************************************************
Enter choice:
```

e) Select menu option 3 to disable the Universal Permanent License.

f) Select option 0 to exit the manage_plr utility.
g) Type **exit** to exit the shell.
h) If applicable, exit the command line interface:

   Versions 6.3 and 6.4: If the CLI is enabled, enter the **exit** command.

   Version 6.5 and later: Enter the **exit** command.

What to do next

(Optional) If you wish to use the system with standard, connected Smart Licensing, see information about Smart Licensing in the Secure Firewall Management Center online help.

Troubleshooting

How do I identify a particular Firepower Management Center in the Product Instance list in Cisco Smart Software Manager?

On the **Product Instances** page in Cisco Smart Software Manager, if you cannot identify the product instance based on a value in one of the columns in the table, you must click the name of each generic product instance of type **FP** to view the product instance details page. The **UUID** value on this page uniquely identifies one Firepower Management Center.

In the Firepower Management Center web interface, the UUID for a Firepower Management Center is the **Product Instance** value displayed on the **System > Licenses > Universal Licenses** page.

I have enabled Universal Permanent Licensing, but now I do not see a Smart License page.

This is the expected behavior. When you enable Universal Permanent Licensing, Smart Licensing is disabled. You can use the Universal Permanent License page to perform licensing operations.

If you want to use Smart Licensing, you must return the Universal Permanent License. For more information, see Deactivate and Return a Universal Permanent License, on page 309.

I was interrupted in the middle of the licensing process. How can I pick up where I left off?

If you have generated but not yet downloaded an Authorization code from Cisco Smart Software Manager, you can go to the Product Instance page in Cisco Smart Software Manager, click the product instance, then click **Download Reservation Authorization Code**.

I have disabled Universal Permanent Licensing, but forgot to copy the Return Code. What should I do?

The Return Code is saved in Firepower Management Center. You must re-enable the Universal Permanent License from the shell (see Enable the Permanent Licensing Menu Option, on page 305.) Then refresh the Firepower Management Center web interface. Your Return Code will be displayed.

What do I do in case of a re-image?

Return the Permanent License, add the Return Code in Cisco Smart Software Manager to release the Permanent License, and re-image the Firepower Management Center box.

What do I do in case of a Return Materials Authorization (RMA)?

Contact Cisco TAC to release the reserved licenses.

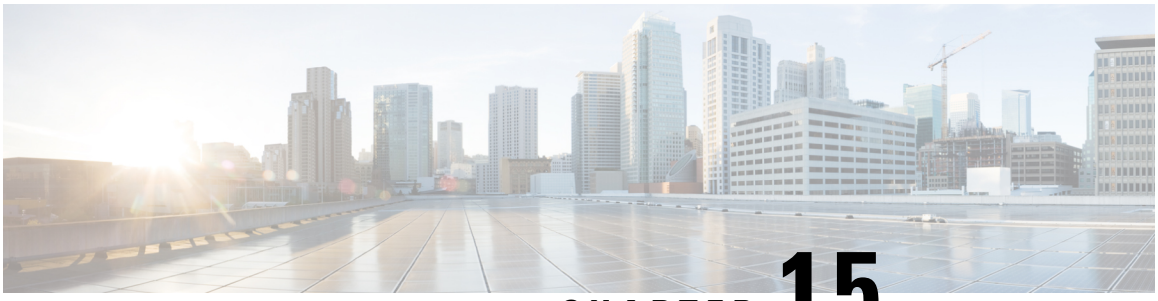
For Assistance

If you have any questions about installing or using Firepower feature licenses, contact TAC Support:

- Visit the Cisco Support site at http://support.cisco.com/.
- Email Cisco Support at mailto:tac@cisco.com.
- Call Cisco Support at 1.408.526.7209 or 1.800.553.2447.

.

Copyright © 2018-2020



CHAPTER 15

Security certifications compliance

These topics describe how to configure your system to comply with security certifications standards:

- Security certifications compliance modes, on page 315
- Security certifications compliance characteristics, on page 316
- Security certifications compliance recommendations, on page 317

Security certifications compliance modes

A security certifications compliance mode refers to a set of security standards established by the U.S. Department of Defense and global certification organizations. Organizations must use only equipment and software that comply with these standards.

Security certification standards

Secure Firewall supports compliance with these security certifications standards:

- **Common Criteria (CC)**: A global standard established by the international Common Criteria Recognition Arrangement, defining properties for security products.
- **Unified Capabilities Approved Products List (UCAPL)**: A list of products meeting security requirements established by the U.S. Defense Information Systems Agency (DISA).

Note The U.S. Government has changed the name of the Unified Capabilities Approved Products List (UCAPL) to the Department of Defense Information Network Approved Products List (DODIN APL). References to UCAPL in this documentation and the Cloud-Delivered Firewall Management Center web interface can be interpreted as references to DODIN APL.

- **Federal Information Processing Standards (FIPS) 140**: A requirements specification for encryption modules used in secure systems.

Organizations can enable security certifications compliance in CC mode or UCAPL mode. Enabling these modes does not guarantee compliance with every requirement of the selected security mode. For more information on hardening procedures, refer to the product guidelines provided by the certifying entity.

Caution After you enable a security certifications compliance mode, you cannot disable it. To remove an appliance from CC or UCAPL mode, you must reimage the appliance.

Security certifications compliance characteristics

This table describes changes to system behaviors and compliance characteristics when CC (Common Criteria) or UCAPL (Unified Capabilities Approved Product List) mode is enabled. Restrictions on login accounts apply to command line access. These restrictions do not apply to web interface access.

System Change	Cloud-Delivered Firewall Management Center		Classic Managed Devices		Secure Firewall Threat Defense	
	CC Mode	UCAPL Mode	CC Mode	UCAPL Mode	CC Mode	UCAPL Mode
FIPS compliance is enabled.	Yes	Yes	Yes	Yes	Yes	Yes
The system does not allow remote storage for backups or reports.	Yes	Yes	—	—	—	—
The system starts an additional system audit daemon.	No	Yes	No	Yes	No	No
The system boot loader is secured.	No	Yes	No	Yes	No	No
The system applies additional security to login accounts.	No	Yes	No	Yes	No	No
The system disables the reboot key sequence Ctrl+Alt+Del.	No	Yes	No	Yes	No	No
The system enforces a maximum of ten simultaneous login sessions.	No	Yes	No	Yes	No	No
Passwords must be at least 15 characters long, and must consist of alphanumeric characters of mixed case and must include at least one numeric character.	No	Yes	No	Yes	No	No
The minimum required password length for the local `admin` user can be configured using the local device CLI.	No	No	No	No	Yes	Yes
Passwords cannot be a word that appears in a dictionary or include consecutive repeating characters.	No	Yes	No	Yes	No	No
The system locks out users other than `admin` after three failed login attempts in a row. In this case, the password must be reset by an administrator.	No	Yes	No	Yes	No	No
The system stores password history by default.	No	Yes	No	Yes	No	No

System Change	Cloud-Delivered Firewall Management Center		Classic Managed Devices		Secure Firewall Threat Defense	
	CC Mode	UCAPL Mode	CC Mode	UCAPL Mode	CC Mode	UCAPL Mode
The `admin` user can be locked out after a maximum number of failed login attempts configurable through the web interface.	Yes	Yes	Yes	Yes	—	—
The `admin` user can be locked out after a maximum number of failed login attempts configurable through the local appliance CLI.	No	No	Yes, regardless of security certifications compliance enablement.	Yes, regardless of security certifications compliance enablement.	Yes	Yes
The system automatically rekeys an SSH session with an appliance: • After a key has been in use for one hour of session activity • After a key has been used to transmit 1 GB of data over the connection	Yes	Yes	Yes	Yes	Yes	Yes
The system performs a file system integrity check (FSIC) at boot-time. If the FSIC fails, Secure Firewall software does not start, remote SSH access is disabled, and you can access the appliance only via local console. If this happens, contact Cisco TAC.	Yes	Yes	Yes	Yes	Yes	Yes

Security certifications compliance recommendations

Cisco recommends that you observe the following best practices when using a system with security certifications compliance enabled:

- Enable security certifications compliance first on the Cloud-Delivered Firewall Management Center, then enable it in the same mode on all managed devices.
- The Cloud-Delivered Firewall Management Center will not receive event data from a managed device unless both are operating in the same security certifications compliance mode.
- For all users, enable password strength checking and set the minimum password length to the value required by the certifying agency.
- To use Cloud-Delivered Firewall Management Centers in a high-availability configuration, configure them both to use the same security certifications compliance mode before forming the high availability pair.

- When you configure Firewall Threat Defense on a Firepower 4100/9300 to operate in CC or UCAPL mode, you should also configure the Firepower 4100/9300 to operate in CC mode. For more information, see the *Cisco Firepower 4100/9300 FXOS Chassis Manager Configuration Guide*.

Do not configure the system to use any features:

- Email reports, alerts, or data pruning notifications.
- Nmap Scan, Cisco IOS Null Route, Set Attribute Value, or ISE EPS remediations.
- Third-party client access to the system database.
- External notifications or alerts transmitted via email (SMTP), SNMP trap, or syslog.
- Audit log messages transmitted to an HTTP server or to a syslog server without using SSL certificates to secure the channel between the appliance and the server.

Configure authentication and access settings appropriately for your security certifications compliance mode.

- Do not enable external authentication using LDAP or RADIUS in deployments using CC mode.
- Do not enable CACs in deployments using CC mode.
- Disable access to the Cloud-Delivered Firewall Management Center and managed devices via the Secure Firewall REST API in deployments using CC or UCAPL mode.
- Enable CACs in deployments using UCAPL mode.
- Do not configure SSO in deployments using CC mode.

The system does not support CC or UCAPL mode for:

- Firewall Threat Defense devices in clusters
- Firewall Threat Defense container instances on the Firepower 4100/9300
- Exporting event data to an external client using eStreamer.

Appliance hardening

To further harden your system, see the latest versions of the *Cisco Secure Firewall Management Center Hardening Guide* and the *Cisco Secure Firewall Threat Defense Hardening Guide*.

Refer to these topics for configuration steps:

Identity and access management

- Users, on page 243

Monitoring, Auditing, and Compliance

- Security certifications compliance, on page 315

System and Network Configuration

- Licenses, on page 275
- *SNMP* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

- *Site-to-Site VPNs for Firewall Threat Defense* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *SMTP Server* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Configure SNMP for FXOS* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Configure Dynamic DNS* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Remote Access VPN* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *FlexConfig Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

Alerting and Device-Level Security

- Create an email alert response, on page 493
- Configure email alerting for intrusion events, on page 503
- Create an SNMP alert response, on page 489
- *Configure NTP Time Synchronization for Threat Defense* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *About Configuring Syslog* in the *Cisco Secure Firewall Management Center Device Configuration Guide*

Protect your network

This section provides links to features that you can configure to enhance and protect your network security. Refer to each feature's documentation for detailed configuration guidance.

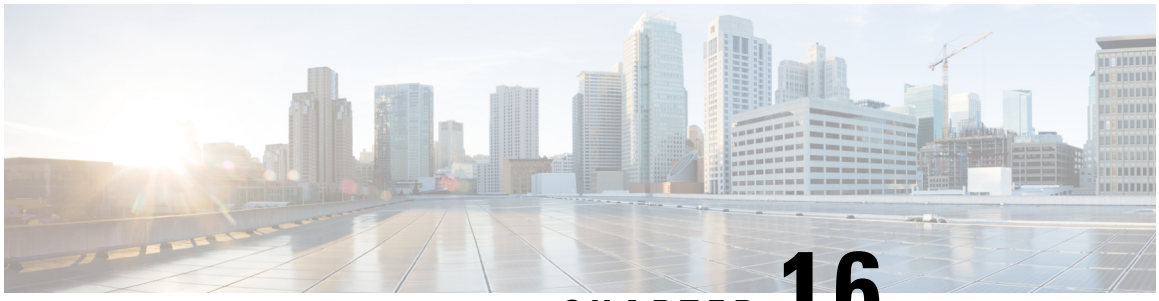
- *Access Control Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*.
- *Security Intelligence* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Getting Started with Intrusion Policies* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Tuning Intrusion Policies Using Rules* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Custom Intrusion Rules* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- Intrusion rule updates, on page 268
- *Transport and Network Layer Preprocessors* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Specific Threat Detection* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Application Layer Preprocessors* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- *Device Management* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
- Updates, on page 263



PART IV

Optimize Firewall Performance with AgenticOps

- Introduction to AgenticOps insights, on page 323



CHAPTER 16

Introduction to AgenticOps insights

AgenticOps brings AI-driven intelligence to firewall management, helping you proactively identify issues, optimize configurations, plan software upgrades, analyze policy health, track feature adoption, and assess capacity requirements. Use AgenticOps capabilities to gain operational visibility, prioritize remediation efforts, and improve the performance, security, and reliability of your firewall deployment.

- About AgenticOps insights, on page 323
- About capacity and traffic insights , on page 331
- About policy change impact insights, on page 334
- About best practices and recommendations , on page 339
- About feature adoption insights, on page 342
- About software upgrade insights, on page 343
- About application insights, on page 345
- About compliance insights, on page 346
- **Manage insight preferences**, on page 349
- Frequently asked questions about AgenticOps, on page 354
- Additional resources, on page 355

About AgenticOps insights

Firewalls are a critical component of any organization's network security architecture. As organizations expand and the threat landscape evolves, managing these firewalls becomes complex. Organizations must continuously update rules and configurations to adapt to new threats, network changes, and compliance requirements, which presents significant challenges. Improper management can lead to security gaps and vulnerabilities. These issues pose risks to an organization's network security.

To effectively address these challenges, a new approach to firewall management is required. This is where AgenticOps becomes essential. AgenticOps leverages artificial intelligence (AI) and machine learning (ML) to streamline and enhance firewall management and network security.

Note Currently, AgenticOps supports:

- Firewall Threat Defense devices that are managed by Cloud-Delivered Firewall Management Center.
- Firewall Threat Defense devices managed by On-Premises Firewall Management Center integrated with Cisco Security Cloud.

  For information on supported versions, see AgenticOps Support for On-Premises Firewall Management Centers.

Key functionalities

- **Traffic and capacity monitoring**: Monitors network traffic and system capacity in real-time, and detects anomalies such as elephant flows. This ensures resources are optimized for peak performance.
- **Policy anomaly detection**: Analyzes firewall policies and detects misconfigurations or anomalies before they impact performance or security.
- **Feature adoption insights and best practice recommendations**: Provides insights into the level of feature adoption and suggests best practices to optimize security configurations.
- **Predictive forecasting for network issues**: Predicts potential future network issues, allowing you to address them proactively and minimize downtime.
- **Operational insights:** Evaluates device readiness for software upgrades, suggests compatible versions, and helps plan upgrades to maintain consistency, stability, and compliance across deployments.
- **Critical alerts**: Filters and prioritizes the most urgent security events, helping you focus on critical issues.

Key features

- **Policy Analyzer and Optimizer**: Analyzes security policies, detects anomalies, and provides recommendations on remediations that can be performed to optimize the policies, thereby improving the firewall performance.
- **Policy Change Impact Analyzer**: Provides visibility into pending access control policy rule changes across devices, analyzes affected rules and traffic segments, and assesses potential traffic impacts to help you make informed deployment decisions.
- **Best Practices and Recommendations**: Generates detailed assessment reports that highlight failed checks against Cisco Secure Firewall best practices and provides actionable recommendations to resolve issues, ensuring optimal firewall performance.
- **Feature Adoption**: Provides insights into the features that are adopted and the percentage of adoption to modify the usage pattern and achieve optimal security. Analyze the adoption rate of different features to improve usage patterns and enhance security measures.
- **Software Upgrade Planner**: Provides upgrade suggestions for your devices through a centralized dashboard. The dashboard displays the current and suggested versions and details about security vulnerabilities and bug fixes.

- **Application Outage Insights**: Detects application outages that may impact users, sites, and business-critical services. The feature correlates outage telemetry, operational insights, and affected resources to help you reduce troubleshooting time.
- **Capacity and Trend Analyzer**: Analyzes firewall capacity utilization and traffic trends to identify abnormal spikes, sustained resource consumption, and emerging performance risks. The feature provides operational visibility into throughput, connection behavior, and traffic patterns to optimize firewall performance.
- **Compliance Posture**: Assess firewall compliance posture against supported security standards. The feature evaluates firewall configurations, identifies compliance gaps, and provides actionable remediation insights to improve security posture
- **Configuration Settings**: Provides the ability to configure thresholds for AgenticOps features and enable or disable insight preferences. You can customize these settings to suit your specific needs.

AgenticOps licensing requirements

If you have licenses for the Secure Firewall Management Center, you can access AgenticOps by enabling AgenticOps Insights in your organization. The initial version of AgenticOps is included as part of your firewall license and is granted on a per-device basis.

Prerequisites to use AgenticOps

- Ensure that AgenticOps Insights is enabled for your Security Cloud Control organization and that Cloud-Delivered Firewall Management Center is provisioned.
- You must have **Super Admin** or **Admin** user roles to opt in and opt out of AgenticOps Insights for your organization.
- If you are using an On-Premises Firewall Management Center, ensure it is integrated with Cisco Security Cloud before onboarding AgenticOps.

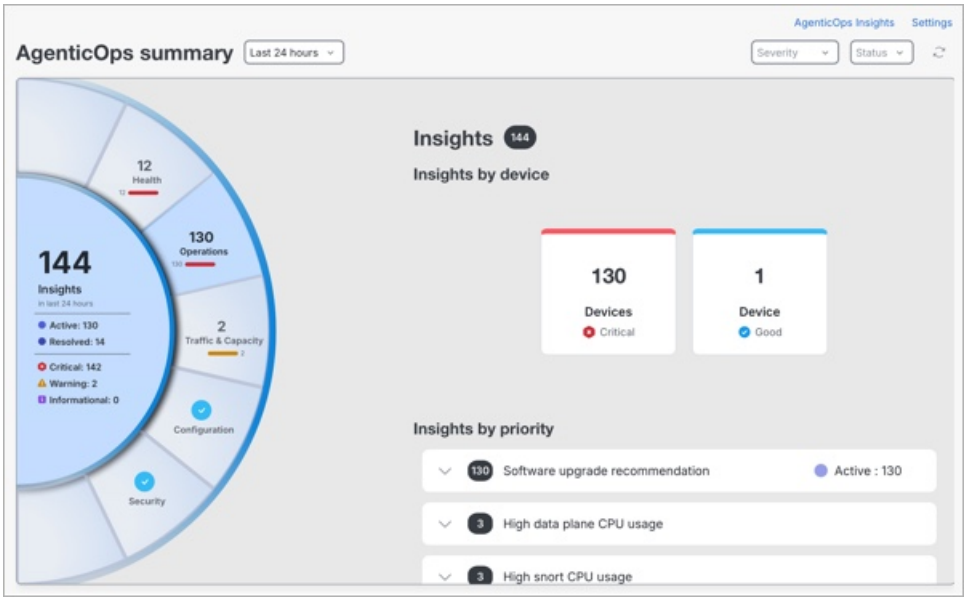
  For information on supported on-prem Firewall Management Center software versions, see AgenticOps Support for On-Premises Firewall Management Center.
- Configure the required thresholds and preferences for supported AgenticOps features.

View summary dashboard

AgenticOps summary dashboard

The **Summary** dashboard provides a consolidated view of all insights across your environment. It enables identification of areas that require attention and gives you the ability to drill down for operational, configuration, and security analysis. You can filter insights by time range, severity, and status.

Figure 34: AgenticOps summary



- **Insights visualization panel**: Visual representation of your insights, helping you quickly understand overall system posture and where to prioritize action.
  - The inner ring summarizes insights by status and severity.
  - The outer ring represents insight categories. Selecting a category displays related details.

    The insight categories include:

    - **Operations**: Focuses on insights that help maintain the operational health and software lifecycle readiness of your devices. This includes upgrade recommendations for your Firewall Threat Defense devices and End-of-Life notifications.
    - **Configuration**: Focuses on insights that evaluate your network policies, configurations, and adherence to Cisco Firewall best practices. This includes identifying misconfigurations, policy anomalies, and deviations that may impact performance, security, or compliance.
    - **Security**: Focuses on insights that detect suspicious, risky, or anomalous activities within your environment. These insights highlight behaviors that may indicate compromised accounts and malicious intent.

  Select any insight category to view the insights associated with it.
- **Insights by device**: Click on a device to view insights.
- **Insights by priority**: Expand the section to view all devices with related insights.
- Use the icons at the top right of the page for additional actions:
  - **AgenticOps Insights**: Navigate to the view of all AgenticOps insights.
  - **Settings**: Navigate to configure preferences and thresholds for insights.

Confidence levels in AgenticOps

A confidence level represents the reliability of a forecast-based insight generated by AgenticOps. It indicates how accurately AgenticOps can predict future behavior for a given condition, helping you decide when immediate action is required versus when monitoring is sufficient.

Applicable forecast-based insight types

Confidence levels are displayed for the following forecast-based AgenticOps insights:

- Remote Access VPN – Maximum Sessions Nearing Limit
- CPU (LINA) Utilization Forecast Breach
- CPU (Snort) Utilization Forecast Breach
- Memory (LINA) Utilization Forecast Breach
- Memory (Snort) Utilization Forecast Breach

How AgenticOps determines confidence levels ?

For forecast-based insights, such as **High Snort Memory Usage**, **High Snort CPU Usage**, and **RA VPN Forecast,** AgenticOpsevaluates how accurately its machine learning models can predict future behavior using historical patterns and current trends.

Based on this forecast accuracy, each insight is automatically assigned a confidence level. Insights that do not meet a minimum reliability threshold are not generated.

- **High**
  - The forecast is highly reliable and strongly supported by historical data.
  - Recommended action: Prioritize the investigation and take proactive corrective measures.
- **Medium**
  - The forecast is reliable but may show moderate variability.
  - Indicates a strong trend that is likely to continue.
  - Recommended action: Review the insight and plan mitigation.
- **Low**
  - The forecast indicates a possible emerging condition.
  - Limited data or higher variability reduces certainty.
  - Recommended action: Monitor closely and validate with additional operational context.
- **Very low**
  - The forecast signal is weak or in an early stage.
  - Higher uncertainty due to insufficient or inconsistent data.
  - Recommended action: Observe trends. No immediate action is required.

Insight statuses and transitions

This table outlines the possible insight statuses, their descriptions, transitions, and examples.

Status	Description	Transition	Triggered by	Example
Active	• An issue is detected and ongoing. • This is the initial state when an issue is identified.	–	System	Upgrade options suggested for Firewall Threat Defense devices.
Resolved	• AgenticOps automatically marks an active insight as resolved when the issue no longer exists. • Historical details are available for reference.	**Active** to **Resolved**: After you fix the issue and the system confirms it in the next check.	• Automatically by the system for performance issues. • By the user (with system confirmation) for configuration issues.	Overlapping firewall rules corrected by the user.

Status	Description	Transition	Triggered by	Example
Not Applicable (N/A)	• The issue existed earlier but it is no longer present. • No historical data is available for reference.	**Active** to **Not Applicable**	System	• When an Access Control policy is deleted, the corresponding insight is marked as NA. • When a device is deleted from FMC inventory, all insights for that device are marked as NA. • When earlier Policy Analyzer and Optimizer results showed issues but a recheck finds zero issues and no details to display, the insight is marked as NA.

View AgenticOps insights

The **AgenticOps Insights** page provides AI-driven alerts that help you detect, prioritize, and resolve issues across your environment.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Summary**.

Step 2 On the **Summary** page, click **AgenticOps Insights** at the top right.

Insights are displayed in chronological order and can be filtered by the following:

- **Time**: Select a time range to view insights generated within that window
- **Severity**: Select a severity level, such as **Critical**, **Warning**, or **Informational**.
- **Insight**: Select a type of insight to view.
- **Category**: Select a category such as **Configuration**, **Health**, **Operations**, and so on.

- **Impacted Resources**: Select a device or service affected by the insight.
- **Status**: Select a status, such as **Active**, **Not applicable**, or **Resolved** stauses.

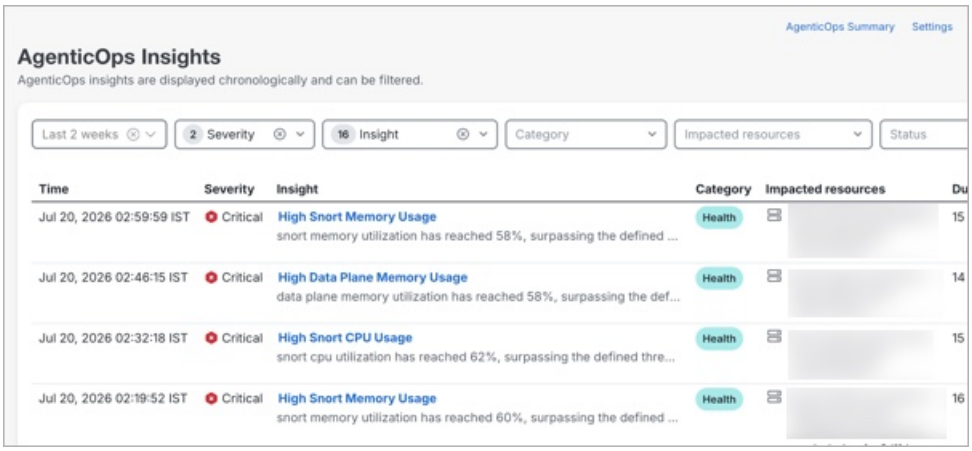
Clicking **Reset all** clears all applied filters and returns to the default view.

Step 3 In the insights table, click an **Insight** or **Impacted Resource** to view additional details.

The details include a summary of the issue, the probable cause, related metrics, and recommended remediation or upgrade actions to ensure optimal performance.

Step 4 Use the icons at the top right of the page for additional actions:

Figure 35: AgenticOps insights



- **AgenticOps Summary**: Navigate to the view the high-level summary of all insights.
- **Settings**: Navigate to configure preferences for AgenticOps features.

AgenticOps for On-Premises Firewall Management Center

AgenticOps supports Firewall Threat Defense devices managed by an On-Premises Firewall Management Center that is integrated with Cisco Security Cloud and onboarded to AgenticOps. Support for individual AgenticOps features depends on the On-Premises Firewall Management Center software version.

Note Currently, On-Premises Firewall Management Center registered through Secure Device Connector (SDC) are not supported.

AgenticOps feature	Supported On-Premises Firewall Management Center versions
Capacity and Trend Analyzer	10.0.0+
Policy Analyzer and Optimizer	7.4.0+

AgenticOps feature	Supported On-Premises Firewall Management Center versions
Software Upgrade Planner	7.4.0+
Best Practices and Recommendations	10.0.0+
Feature Adoption	10.0.0+
Application Insights (ThousandEyes)	7.6.0+
User Risks Insights	10.0.0+
Policy Analyzer and Optimizer - Network Address Translation	7.6 +
Policy Analyzer and Optimizer - Access Control Policy	7.2.0 +

After onboarding, supported On-Premises Firewall Management Center- managed devices are available throughout AgenticOps, including the Summary page, AgenticOps Insights, and supported feature dashboards.

About capacity and traffic insights

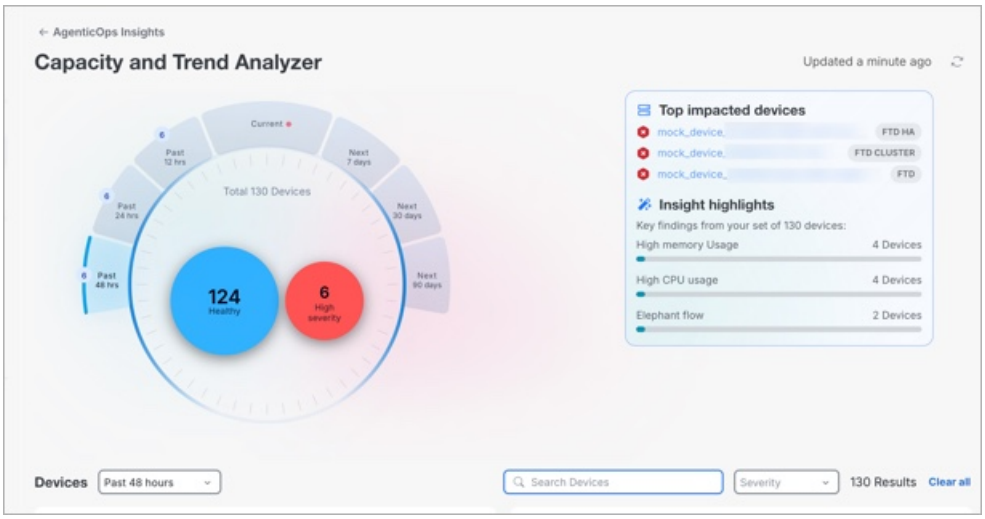
Capacity issues often develop gradually and only become visible when they begin to impact performance. The **Capacity and Trend Analyzer** provides visibility into resource utilization across your devices, highlights emerging risks, and predicts when those risks may lead to performance issues. By combining historical trends with predictive insights, it enables you to anticipate problems and take action before they affect users. This feature enables proactive capacity planning by identifying future risks before they impact system performance.

How AgenticOps generates capacity insights

Forecasted insights predict future resource utilization based on historical trends and patterns. AgenticOps continuously analyzes historical data and runs periodic forecasts to identify when resources are likely to reach configured thresholds.

AgenticOps uses multiple analytical approaches to identify and predict issues. AgenticOps runs forecasts on a recurring weekly basis to identify potential capacity risks in advance, including:

- Detect threshold breaches: Identifies when CPU or memory usage exceeds configured limits in real time.
- Detect anomalies: Identifies unusual deviations from expected behavior based on historical patterns.
- Detect elephant flows: Detects large traffic flows that can impact system performance and resource utilization.
- Detect Remote Access VPN session limits: Detects or predicts when VPN session capacity limits are reached.
- Forecast capacity risks: Predicts future threshold breaches using historical trends, enabling proactive planning.



Note

- The accuracy of insights depends on the availability and quality of historical data. Limited or inconsistent data may delay or reduce insight accuracy.
- In environments with limited activity or newly onboarded devices, insights and forecasts may take time to appear.

Table 17: Key features of Capacity and Trend Analyzer

Feature	Description
Capacity insights dashboard	Provides a fleet-level view of device capacity status, displaying past, current, and predicted capacity insights across all monitored devices.
Capacity insights	Generates insights based on CPU, memory, connection, throughput, and VPN session activity to highlight potential resource constraints.
Device capacity analysis	Displays detailed capacity insights for an individual device, including resource metrics, AI-generated summaries, and forecast information.
Predictive capacity forecasting	Uses historical resource utilization data to predict future usage trends and identify potential capacity risks before thresholds are reached.
Timeline-based analysis	Allows you to analyze capacity insights across different time windows, including past events, current conditions, and predicted future risks.
Device health indicators	Displays device health status based on active and predicted capacity insights, helping you quickly identify devices that require attention.

Supported devices

The Capacity and Trend Analyzer supports:

- Secure Firewall Threat Defense devices
- Secure Firewall Threat Defense High Availability (HA) deployments
- Secure Firewall Threat Defense cluster deployments

Note For High Availability (HA) and cluster deployments:

- Capacity analysis runs at the system level.
- Individual member devices are not displayed separately.

Analyze capacity trends and forecast resource usage

Before you begin

Ensure that **Capacity and Trend Analyzer** is enabled under **Settings**. For more information, see enable Capacity and Trend Analyzer.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Capacity and Trend Analyzer**.

This opens the dashboard where you can view capacity insights across your environment.

Step 2 Use the circular timeline to select a timeframe, such as **Current**, **Past (12 hours**, **24 hours**, or **48 hours)**, **Next (7 days**, **30 days**, or **90 days).**

- **Current** shows active capacity insights.
- **Past 12, 24, or 48 hours** show insights from the selected past period.
- **Next 7, 30, or 90 days** show predicted capacity risks.

The timeline also displays the total number of devices and their health status, such as **Healthy**, **Needs attention**, and **High severity**.

Step 3 Review the summary panel.

- Use the **Top impacted devices** section to identify devices with the most critical capacity issues.
- Use the **Insight highlights** section to understand the most common capacity risks across your organization.

Step 4 Review device insights in the **Devices** section.

- Search for a device or use filters such as timeframe and severity to locate devices that require attention.

- Each device displays a summary of detected or predicted issues, along with associated insight tags.

Step 5 Select a device to view detailed insights.

This view provides summaries and detailed insights for the selected device.

Step 6 On the device details page, review the **AI summary**, **Resolved insights**, **Active insights**, and **Future insights**.

Step 7 In the **Device metrics** section, select a metric to view historical trends and forecast data.

Step 8 Use the insights and forecasts to determine appropriate actions, such as redistributing traffic, scaling resources, or planning capacity upgrades. After taking action, use the **Current** view to monitor whether the issue has been resolved.

The **Capacity and Trend Analyzer** helps you move from reactive troubleshooting to proactive capacity management by enabling you to identify risks early and take informed action.

About policy change impact insights

Changes to an access control policy can affect network traffic in unexpected ways. A modification that appears minor, such as updating a network object, expanding a rule, or changing rule conditions, can alter how traffic is evaluated by other rules in the policy. These effects are often difficult to identify before deployment.

Policy Change Impact Analyzer helps you review the potential impact of pending access control policy rule changes across devices before deployment. The analysis helps you identify affected rules, impacted connections, and traffic segments so that you can better understand the consequences of a change before deploying it. The **Policy Change Impact Analyzer** evaluates proposed policy changes and provides visibility into:

- The overall analysis result for each device
- Affected devices and rules
- Impacted traffic segments
- Historical connections affected by a changed rule, when traffic data is available
- Rule-level inferred impact
- AI-generated summaries that explain the potential effect of a rule change

Limitations and considerations

- Currently, **Policy Change Impact Analyzer** supports Cloud-Delivered Firewall Management Center-managed Firewall Threat Defense devices only.
- **Policy Change Impact Analyzer** is available only for deployments in the US, AUS and EU regions.
- High Availability and cluster deployments are represented by their parent target during analysis.
- Impacted connection counts are available only when **Security Analytics and Logging** is enabled and connection events are sent for FTD Management Services.
- Running an analysis does not deploy, edit, approve, or reject policy changes.

- AI-generated explanations support decision-making. Validate important findings against your security requirements and operational policies before deployment.
- Review analysis results before approving production changes, especially high-risk policy modifications.

Understand analysis results

Review the assessment result to understand the likely effect of pending policy changes and the recommended next steps.



Note An analysis result does not indicate whether a change is approved. Review the transition direction, affected rule, and traffic conditions before deployment.

Analysis result	What it means	Recommended action
Critical	One or more changes may expose the network to threats or disrupt services.	• Review affected rules and traffic segments before deployment. • Revise and rerun when the change is not intended.
Potential impact	One or more changes expand access or may affect traffic.	Validate the affected traffic and the expected business outcome before deploying.
Improved security posture	One or more changes reduce exposure or block risky traffic.	Confirm that the affected traffic is intentionally restricted, then continue through normal change review.
No traffic impact	No behavior-changing traffic transition was identified from the policy comparison.	• Confirm that the proposed rule is needed and is positioned as intended. • This result is different from a change that has zero observed connections.
Indeterminate	The available evidence was insufficient to determine a reliable impact.	• Review the changed rule and traffic criteria manually. • Consider gathering additional operational context and rerunning the analysis.

Understand job statuses

Job Status shows the current state of the policy impact analysis and whether the latest result is available and up to date.

Job status	Description
Verifying analysis is up to date	The system is checking whether the latest analysis still matches the pending policy changes.
Queued	The analysis request was accepted and is waiting to run. Wait for the status to change to **In progress**.
Completed	The latest analysis completed successfully. A warning icon may appear if a policy analysis returned an error.
Partial success	Useful assessment results are available, but one or more stages did not complete. Review the available details and **Workflows** before deciding whether to rerun.
In progress	The system is preparing policy data, calculating impact, or collecting connection evidence. Wait for the result or open **Workflows** to monitor progress.
Out of date	The policy changed after the analysis started, or the analysis is more than 24 hours old and its traffic evidence may no longer be current. Rerun the analysis to get the latest results.
Failed	The latest overall analysis failed. Review the error and **Workflows** details, then rerun the analysis if appropriate.
Cancelled	The analysis was cancelled by an authorized user before completion.
Not analyzed	No analysis is available for the device, or the device has pending changes that have not yet been analyzed. Select a device and run the analysis.
Not applicable	No pending policy changes are available for analysis. Create or save the intended policy changes, then run the analysis.
Unknown	The system could not determine a reliable current status. Refresh the device list and try again later. If the issue persists, review the workflow details.

Understand inferred impact

Inferred Impact describes the likely effect of an individual changed rule or affected rule.

Inferred impact	Description
Potential threat exposure risk	The change may allow risky applications, risky URL categories, broad application or URL scope, sensitive destination ports, or broad Internet exposure.
Potential service disruption	The change may block sensitive services or high-volume business traffic.
Access expansion	The change may allow traffic that was previously blocked without a clear risky-application or sensitive-port signal.
Indeterminate	The available evidence was insufficient, or the impact assessment did not complete.
Improved security	The change blocks risky or untrusted traffic, or reduces broad external exposure.
No traffic impact	No behavior-changing traffic transition was detected.

Impacted traffic segments and connections

An **Impacted Traffic Segment** is a distinct combination of traffic conditions whose policy outcome, such as allow or block, changes when the proposed policy changes are deployed.

- Traffic conditions may include source and destination zones, networks, applications, URLs, users, services, ports, security-group tags, and geolocation identities.
- The number of impacted traffic segments is not the number of rules, packets, IP addresses, or historical connections.

Impacted Connections shows the number of historical connections associated with traffic affected by a changed rule.

- The count is based on available Security Analytics and Logging connection-event data. It represents observed historical activity, not a forecast of future traffic.
- If connection events are unavailable, the analyzer can still identify impacted traffic segments, but historical connection counts are not available.
- Review the policy comparison and traffic segments instead of treating unavailable data as zero traffic.

Understand policy analysis limitations

Some policy constructs are not fully evaluated during analysis. Review affected rules manually before deployment when these constructs are involved.

- Object overrides are not considered during analysis. For example, if a port object has the value `8080` but is overridden to `80` for a particular device, the device-specific value `80` is not considered.
- Country, continent, and custom geolocation objects are expanded to country-level identities. They are not expanded to IP-level identities, and individual IP addresses are not compared.

- URL category reputation is not considered. For example, if a new rule uses the same URL category with a different reputation, the analyzer may report **No Traffic Impact** even though the reputation changed.
- URL list and URL feed objects are not expanded into individual URLs. The analyzer reports an impact only when the object is used in an affected traffic fragment or segment. Individual entries in the list or feed are not evaluated separately.
- Time-range support is limited. Only active time ranges are considered. Rules with future or expired time ranges are treated as disabled. Changes to the time-range configuration itself are not considered during analysis.
- Dynamic-attribute support is limited to Security Group Tags (SGTs), including supported FMC and ISE SGT references. CSDAC-backed dynamic attributes and other unsupported dynamic-object criteria are not analyzed.
- Application filters entered directly in the application field, such as **Very High** or **High** risk or business relevance criteria, are not supported when they are inline filters rather than user-defined application filter objects. User-defined application filters, including filters that combine multiple criteria, are supported.

Analyze policy change impact

Before you begin

- Enable the **Policy Change Impact Analyzer** toggle under **Insights & Reports** > **Settings**.
- To view historical impacted connection counts, enable **Security Analytics and Logging**.
- **Policy Change Impact Analyzer** currently supports Cloud-Delivered Firewall Management Center-managed Firewall Threat Defense devices.

Procedure

Step 1 In the left pane, click **Insights & Reports > Policy Change Impact Analyzer**.

The **Policy Change Impact Analyzer** page displays an **Impact Analysis** summary, including the number of identified impacts.

The device list displays:

- **Device Name**
- **Job Status**
- **Analysis Result**

Step 2 Select one or more devices, and then click **Run Impact Analyzer**.

You can select up to 10 devices for analysis. The analyzer evaluates the pending policy changes for the selected devices.

Step 3 Monitor the analysis in the device list.

Use the search field to search by device name, analysis result, or job status.

Step 4 Click a device row to open the details pane.

The pane displays the device status and may include these actions:

- **View change impact details** to open the detailed analysis.
- **Re-analyze** to generate an updated result.
- **Deploy** to open the standard deployment workflow.
- **Workflows** to review analysis execution details and errors.

Step 5 Click the device name to open the analysis report.

The analysis details page displays the overall impact assessment for the selected device.

Step 6 Review the **Analysis Result**, **Rule Change Summary**, **Impacted Traffic Segments**, **Impacted Connections**, and **Inferred Impact**.

Step 7 If the analysis is incomplete, failed, out of date, or indeterminate, review the policy manually before deployment and rerun the analysis as needed.

Note
Policy Change Impact Analyzer does not deploy, edit, approve, or reject policy changes. Use the standard deployment and change-control process after reviewing the analysis results.

The **Policy Change Impact Analyzer** displays the potential impact of pending access control policy changes, including changed rules, impacted traffic segments, historical connections, inferred impact, and rule change summaries.

About best practices and recommendations

Enhance your organization’s security posture with AgenticOps by identifying deviations from Cisco Secure Firewall best practices. Run assessments on your devices, generate reports, and receive insights that guide you toward optimal performance.

- Assessment: Evaluates your firewall configuration across multiple categories. Each check determines alignment with Cisco Secure Firewall best practices. The report summarizes the total number of checks performed. It categorizes the results as **Passed** or **Requires review**.

  Checks that require review indicate deviations that could impact firewall efficiency and security. Each failed check presents an opportunity for improvement. Addressing these checks contributes directly to optimizing firewall performance.

- Recommendation: Provides specific recommendations to address identified issues, ensuring optimal firewall performance. These include detailed information such as the nature of the problem, symptoms, impact, and required actions.

The best practices and recommendations checks are developed with input from Cisco's Technical Assistance Center (TAC) and Customer Experience (CX) teams. This input helps address trending issues, incorporate industry best practices, and enhance the reliability of recommendations. Implement these recommendations to resolve issues, align with best practices, and optimize firewall performance.

Table 18: Key features

Feature	Description
Automated assessments	Runs periodic evaluations of firewall devices against Cisco best practices.
Checks summary	Displays how many checks passed and highlights those requiring review.
Trend visualization	Shows the number of checks over time, helping you compare passed and failed checks across assessment cycles.
Device reports	Provides device-specific results and percentage of improvement potential.
Review category and check control	Enable or disable review categories or individual checks for future assessments.

Implement best practices and recommendations

Before you begin

Ensure that **Best Practices** is enabled under **Settings**. For more information, see Enable Best Practices.

Procedure

Step 1 In the left pane, click **Insights & Reports > Best Practices and Recommendations**.

The **Assessment Summary** provides a high-level overview of assessment results. It includes two tiles:

- **Checks summary**: Displays the total number of checks and highlights those requiring review.
- **Best practices assessment trend**: Helps you track assessment outcomes over time. The Y-axis represents the number of checks, and the X-axis shows assessment dates. You can hover over data points to view summary statistics.
- **Feedback**: A mechanism for you to provide feedback on the assessments. You can enter optional comments and grant consent for follow-up contact regarding your feedback.

Step 2 In the **Device reports** section, you can view the list of all available device reports. Filter devices by **Device status, Review categories,** or **Assessment status** to narrow down results.

Assessment Statuses: Each device report has an **Assessment Status**, which indicates the current state of the assessment.

- In progress: An assessment is actively running. After completion, a report will be generated.
- In queue: The previous assessment is outdated, and a new one has been scheduled.
- Updated: The assessment is complete, and the latest report is available for review.

- Error: The assessment could not be completed due to an error. The report will be automatically generated after 24 hours. If the issue persists, contact **Cisco TAC** for assistance.

Step 3 From the three-dot menu icon next to each device:

a. Click **Run assessment** to initiate a new assessment.

b. Click **Download report** to export the Best Practices assessment summary report in PDF format.

Schedule Device Assessment

Periodic assessments run automatically, but you can also run assessments manually at any time. In environments with more than **50 devices**, only **50 devices can be analyzed per day**. Use the **Select devices** option on the **Best Practices and Recommendations** page to choose which devices are included in automated assessments.

- **Select devices for automated assessments**:
  a. Click **Select devices** in the notification banner to open the **Schedule devices for assessment** dialog. This dialog allows you to choose which devices run automated assessments.
  b. Devices in the **Scheduled devices** panel are included in automated assessments. Click **Save**.
- **Mange scheduled devices**: You can update the list of scheduled devices at any time. Devices removed from the **Scheduled devices** panel are no longer included in automated assessments. Devices added to this panel are included in future automated assessments.
- **Device scheduling limits**: You can schedule a maximum of **50 devices** for automated assessments. The **Scheduled devices** panel displays the number of currently scheduled devices, for example **48 / 50 devices**. When 50 devices are selected, further selection is disabled and a tooltip indicates that the maximum number of scheduled devices has been reached.

Step 4 Click on a **Device name** to view a detailed report for that specific device.

- Check the **Show passed checks** checkbox to view successful checks in the detailed view. This provides full visibility into all checks.
- In the **Best practices assessment** section, view the **Total checks**, how many **Passed** and **Require review**.
- Expand each check to view the remediations and corrective actions.

Step 5 Enable or disable an entire category from future assessments. You can also disable individual checks within a category.

- Disable the toggle to exclude a review category or check from future assessments.
- Enable the toggle to include it in future assessments.

Note
Disabling a category or check does not affect the actual feature or its operation. The feature continues to function normally, but it is excluded from **Best Practices** assessments.

About feature adoption insights

The **Feature Adoption** dashboard provides a comprehensive view of how effectively different features are being utilized across your environment. Each feature has an adoption rate. This rate is calculated based on the total number of eligible devices and the number of those devices configured with the feature. This insight helps you identify underutilized capabilities and take steps to improve overall adoption.

The dashboard groups features into license categories such as **Essentials**, **IPS**, **Remote Access VPN**, **Features not requiring license**, and **Subscription-Based Features** to help you understand adoption patterns.

Note If a license is not available for a given category, the adoption rate for that category will always be **0%**.

By tracking the adoption rate at both the feature and license category levels, you can prioritize improvements and maximize the usage of available features. **Feature Adoption** helps you use available capabilities to achieve your intended outcomes

Assess and improve feature adoption

Before you begin

Ensure that you have selected the list of features under **Insights & Reports** > **Settings** > **Feature Adoption**.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Feature Adoption**.

- The **Summary** tile provides a quick overview of overall adoption status:
  - **Adoption rate**: Percentage of features currently adopted across their applicable scope.
  - **Total features**: Total number of features available for adoption.
  - **Not adopted**: Number of features that have not been enabled or configured.
  - **Partially adopted**: Number of features that are enabled or in use but not fully across all applicable scope.
  - **Adopted**: Number of features that are fully enabled and in use.
- In the **Feature Recommendation** tile, you can watch short videos about recommended features that will help enhance your organization's security.

Note

- The data updates every 24 hours, but you can click **Refresh** to update it manually.
- We recommend that you increase the usage of these features to improve overall performance.

Step 2 Click on a feature name to view the following details:

- A short description of the feature.
- The feature adoption rate can range from 0% to 100%, based on how much the feature is used. Enable the toggle to include the feature in adoption calculations or disable it to exclude it.

  Note
  Enabling or disabling a feature affects only the adoption score and does not impact the functionality of the feature.

- Steps to improve your feature adoption rate.

About software upgrade insights

The **Software Upgrade Planner** helps simplify firewall software upgrade planning and management. It analyzes your deployed Firewall Threat Defense devices and provides upgrade recommendations based on compatibility, stability, security vulnerabilities, and bug fixes. This helps you make informed upgrade decisions and maintain a consistent software version across the environment.

The feature provides a centralized view of the current software version running on each device along with the suggested version. It also highlights important details such as bug fixes, vulnerability fixes, and upgrade readiness insights. By identifying potential risks and compatibility concerns before the upgrade process begins, the planner helps reduce operational issues and minimizes downtime.

Need a visual guide? Explore the Software Upgrade Planner feature walkthrough to see it in action.

Guided remediation

Software Upgrade Planner also provides guided remediation insights for vulnerabilities and bugs affecting your devices.

When you view the details of a security vulnerability or bug, guided remediation provides AI-generated insights to help you understand the issue and determine an appropriate upgrade version. These insights can include:

- **Root cause analysis**: Summarizes the issue and its potential impact.
- **Workaround**: When available, helps mitigate the issue before upgrading.
- **Remediation guidance**: Identifies how many devices in your deployment are affected and recommends whether upgrading to a specific software version resolves the issue.

Threat campaign insights

Software Upgrade Planner identifies devices that are exposed to active threat campaigns. When applicable, you can view the active threat campaigns affecting a device and access Cisco Talos intelligence resources to learn more about each campaign, including its scope, impact, and recommended mitigations.

Manage software upgrades

Use the Software Upgrade Planner to review suggested software versions for your Firewall Threat Defense devices and evaluate the security and operational benefits of upgrading. In addition to suggested upgrade versions, the planner provides insights into active threat campaigns and guided remediation for vulnerabilities and bugs affecting your devices.

Before you begin

Ensure that the **Software Upgrade Planner** toggle is enabled under **Insights & Reports** > **Settings** > **Operations**.

Procedure

Step 1 In the left pane, click **Insights & Reports > Software Upgrade Planner**.

The **Software Upgrade Planner** page displays a summary of upgrade opportunities across your deployment. The **Summary** section provides an overview of:

- Devices with upgrades available
- Total available fixes
- Details on threat campaigns
- Security vulnerability fixes
- Bug fixes

You can click **View all** in the **CVE Fixes** or **Bug Fixes** tiles to review the corresponding issues across your deployment.

Step 2 Review the suggested upgrade versions for your devices.

- Multiple upgrade options are available based on vulnerabilities, bugs, and new features. The Cisco-suggested version is indicated by a gold star. Choose the version that suits your requirements.
- For each device, the planner displays the current software version and suggested upgrade versions. If a device is affected by one or more active threat campaigns, the current software version displays the number of threat campaigns affecting the device.
- Hover over the **Current version** to view additional details, including the current software version, vulnerabilities found, bugs found, and the active threat campaigns affecting the device.
- Select a threat campaign to open the corresponding Cisco Talos Intelligence blog and learn more about the campaign.

Step 3 To perform an upgrade, click **More actions** (⋮) for the device, and click **Go to product upgrade**

The **Product Upgrades** page opens in Cloud-Delivered Firewall Management Center, where you can perform the upgrade.

- View available upgrade packages.

- Perform an upgrade by clicking **Upgrade** in the **Actions** column. For more information, see Cisco Secure Firewall Threat Defense Upgrade Guide for Cloud-delivered Firewall Management Center.

Step 4 In the **Software Upgrade Planner** page, select a device to view more details.

The device details page displays **Suggested** and **Golden** versions. For each suggested version, you can review:

- Active threat campaigns resolved
- Security vulnerability fixes and bug fixes
- Estimated downtime
- New features included in the release
- Link to detailed release notes

Use this information to compare the available upgrade versions and determine the most appropriate upgrade version for your deployment.

Step 5 Review the **Security vulnerability fixes** or **Bug fixes** tabs.

- These tabs list the vulnerabilities and bugs affecting the selected device. You can search and filter the results or export the list as a CSV file.
- Select a CVE ID or bug ID to open the details pane. The details pane provides AI-generated insights to help you understand the issue, including:
  - **Root cause analysis** that summarizes the issue and its potential impact.
  - **Workaround** information, when available, to help mitigate the issue before upgrading.
  - **Remediation** guidance that identifies how many devices in your deployment are affected and explains whether upgrading resolves the issue.
- In the **Bug fixes** tab, bugs are categorized as **Impacting bugs** and **Other bugs**. Impacting bugs are affecting the selected device and its environment. You can search the list, filter bugs by status (**Fixed bugs**, **Open bugs**, or **All**), and export the results as a CSV file.

You can use the **Software Upgrade Planner** to evaluate suggested software versions, review threat campaign exposure, analyze security vulnerabilities and bug fixes, and use AI-generated remediation insights to make informed upgrade decisions.

About application insights

Application Insights identifies outages for monitored cloud applications by consuming outage intelligence provided by **ThousandEyes Internet Insights.**

ThousandEyes Internet Insights collects telemetry from globally distributed vantage points that continuously measure application and network availability. When an outage affecting a SaaS or cloud service is detected, an outage event is generated. Application Insights uses this intelligence to display outage details such as affected applications, outage duration, impacted regions, and affected domains for the applications you monitor.

For more information about **ThousandEyes Internet Insights**, see Internet Insights.

How application insights work

- Once **Application Insights** is enabled, AgenticOps periodically evaluates Cloud-Delivered Firewall Management Center policies to determine which applications are eligible for monitoring.
- The Cloud-Delivered Firewall Management Center policy is scanned every 24 hours. Applications that are enabled in the policy are automatically monitored, and any detected outages are reported on the **Application Insights** page using outage intelligence from **ThousandEyes Internet Insights**.
- On the **Settings** page, applications that are allowed in the Cloud-Delivered Firewall Management Center policy are preselected by default. You can optionally select additional applications from the list to include them in outage monitoring.

Need a visual guide? Explore the Application Insights feature walkthough to see it in action.

Detect application outages

Before you begin

Ensure that the **Application outage insights** toggle is enabled under **Insights & Reports** > **Settings** > **Application Insights**.

Procedure

Step 1 In the left pane, click **Insights & Reports > Application Insights**.

The Application Insights page displays a summary of detected outages, including the total number of outages, active outages, resolved outages, and the number of applications being monitored for the selected time range.

Step 2 Review the list of affected applications. For each application outage, the following information is displayed:

- Application name
- Affected domains
- Outage duration
- Start and end timestamps
- Affected regions

Step 3 Click an application to view detailed outage information in the right pane.

About compliance insights

Organizations spend significant time and effort validating firewall configurations against industry security standards. This process is often manual, inconsistent, and dependent on specialized expertise. As compliance

requirements evolve, teams must repeatedly revisit configurations, which slows down deployments and increases operational overhead.

Even small misconfigurations such as, overly permissive access rules or missing segmentation controls, can lead to compliance violations that are difficult to detect through manual review.

Compliance Posture in AgenticOps evaluates firewall configurations against industry-standard requirements. It translates compliance guidelines into programmatic checks that can be applied consistently across environments. The current implementation evaluates firewall configurations against the **Payment Card Industry Data Security Standard (PCI DSS) v4.0.1 framework**. Currently, **Compliance Posture** is available only for deployments in the US, EU, and AUS regions.

Disclaimer: The report is provided for informational purposes only and is created using AI. It does not constitute legal or professional compliance advice and is not a substitute for a formal audit or assessment performed by a qualified security assessor. You are responsible for independently verifying the contents of the report for accuracy and completeness. Cisco disclaims all liability related to the report.

Benefits

- **Automate compliance validation** by evaluating firewall policies against predefined rules, eliminating manual review and reducing the risk of human error.
- **Generate compliance reports** that provide detailed findings and recommendations to help you assess and improve your security posture.
- **Review actionable recommendations** that guide remediation, such as restricting overly broad access, validating business intent, or removing redundant rules.
- **Review security zone mappings** to evaluate segmentation controls based on relationships between network zones.

Prepare your environment for compliance evaluation

Before generating a compliance report, make sure that:

- Cloud-Delivered Firewall Management Center is provisioned in your organization.
- AgenticOps is activated in your organization.
- Firewall Threat Defense devices are onboarded and managed through Security Cloud Control Firewall Management.
- The device that you want to evaluate is online and reachable.

Generate compliance reports

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Compliance Posture**.

The **Compliance Posture** page displays previously generated reports and allows you to create new compliance evaluations.

Step 2 Click **Generate Report** to begin defining the report.

Step 3 Enter a name and choose the device to evaluate.

- Currently, **PCI DSS 4.0.1 Compliance Template** is the only supported template. This template evaluates firewall configurations against requirements for network security, access control, secure configuration, and logging.
- Only devices that are currently online are available for selection.

Step 4 Review the security zones.

PCI DSS requires the Cardholder Data Environment (CDE) to be accurately defined. The CDE includes systems that store, process, or transmit cardholder data or sensitive authentication data, as well as environments with unrestricted connectivity to systems handling such data. An automated assessment identifies security zones that are likely to be part of the CDE based on the existing configuration. Review the selected security zones and update the selections based on your knowledge of the network before generating the report.

Step 5 In the **Security Zones** table, select the checkbox in the **CDE** column for each security zone that is part of the CDE. Clear the checkbox for any zone that is not part of the CDE.

Note
Incorrect security zone selections can result in inaccurate compliance findings, particularly for checks related to segmentation and traffic between network zones.

Step 6 Click **Generate Report** to begin the compliance evaluation process.

Note
- Compliance reports are generated manually and are not evaluated automatically.
- The system analyzes firewall policies and validates them against the selected compliance template. Report generation can fail due to device connectivity issues. Verify that the device is online and reachable.
- The page confirms that report generation has started. A comprehensive review and analysis of the policies on the selected device will be performed. This process may take some time.
- Click **Return to compliance reports** to return to the **Compliance Posture** page.

Step 7 Find the generated report.

Use the search field to search reports by name, user, or template. You can also click **Filters** to narrow the results. The **Compliance Status** column shows the number of checks that passed, need review, or failed.

Step 8 Click the report name to view the detailed compliance results.

- The report page displays the report name, evaluated device, last update time, compliance score, and compliance summary.
- The **Compliance Score** shows the number of compliant checks out of the total number of checks.
- The **Compliance Summary** displays total findings, checks that passed, need review, or failed.
- The report displays a list of **Compliance Checks**, each representing a specific requirement being evaluated (for example, restricting inbound traffic to the CDE or validating VPN security). Select a compliance check to view its details, which may include:
  - A description

- References to relevant compliance standards
- Identified findings associated with the check
- Recommendations for remediation

The number next to each compliance check indicates how many findings were identified for that check.

Step 9 Click **Download Report** to export the report in JSON format for audit and offline review.

Step 10 Rerun a report.

After making changes to address compliance findings, rerun the report to evaluate the updated configuration.

You can view compliance findings from the **Summary** page. Navigate to **Insights & Reports > Summary** and click **Operations** to view compliance-related insights for the selected time range.

Manage insight preferences

You can enable or disable insight preferences for your tenant for the following AgenticOps features:

Enable AgenticOps insights

To take advantage of AgenticOps' benefits, you must enable AgenticOps insights. You must have **Super Admin** or **Admin** user roles to enable **AgenticOps Insights** in your tenant.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Settings**.

Step 2 Select **Confirm AgenticOps activation**, and click **Get Started**.

Step 3 Review the available AgenticOps features for each detected Firewall Management Center. The page displays the supported AgenticOps features, software version, and the number of managed devices for each Firewall Management Center. Click **Setup AgenticOps**.

Step 4 Click **Confirm Onboard**.

The onboarding process begins, and it takes a few minutes to fetch the data that is required to provide the insights. When completed, the **AgenticOps** > **summary** page is displayed.

Disable AgenticOps insights

You can choose to opt out of AgenticOps if you no longer want to use the capabilities. When you opt out, AgenticOps stops collecting data from your devices and deactivates all data processing and alerts. You must have **Super Admin** or **Admin** user roles to disable AgenticOps insights in your tenant.

Data cleanup options

When opting out, you can decide how existing data is handled:

- **Keep existing data**: Stops the AgenticOps service but retains your historical data. If you choose to opt in again later, this data can be reused for improved insights.
- **Remove all data**: Stops the AgenticOps service and permanently deletes all existing AgenticOps data. Deleted data cannot be recovered.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Settings**.

Step 2 Click **Miscellaneous**.

Step 3 Under **Opt out of AgenticOps**, choose your preferred data cleanup option.

Step 4 Click **Opt out of AgenticOps**.

Manage AgenticOps data sources

Before you begin

- Ensure that AgenticOps is enabled for your organization.
- If you're using an On-Premises Firewall Management Center, ensure that it is integrated with Cisco Security Cloud.

Manage the data sources such as Cloud-Delivered Firewall Management Center and On-Premises Firewall Management Centers registered with Cisco Security Cloud, that provide information to AgenticOps. You can onboard the data sources to enable AgenticOps capabilities or offboard data sources that you no longer want AgenticOps to use.



Note Cloud-Delivered Firewall Management Center is required for AgenticOps and can't be offboarded.

Onboard a data source

1. In the left pane, click **Insights & Reports > Settings**.
2. Under **AgenticOps Settings**, click **Data Sources**.
3. When AgenticOps detects a supported data source that isn't onboarded, a banner appears on the **Data Sources** page. Click **Review & Onboard**.
4. Review the supported AgenticOps features for the selected data source.
5. Click **Confirm Onboard**.

The selected data source is onboarded and becomes available for supported AgenticOps features.

Offboard a data source

1. In the left pane, click **Insights & Reports > Settings**.
2. Under **AgenticOps Settings**, click **Data Sources**.
3. Locate the onboarded data source.
4. Turn off the toggle for the data source.

AgenticOps no longer collects data or generates insights from the selected data source.

Capacity and traffic insights

You can modify the preferences for Capacity and Traffic-related insights. These insights help monitor firewall performance, detect anomalies, and identify potential capacity risks.

Procedure

Step 1 In the left pane, click **Insights & Reports > Settings**.

Step 2 Click **Capacity & Traffic Analysis**.

Step 3 Enable the toggle buttons for the insights you want to monitor:

a. Traffic-related insights

- High traffic caused by elephant flow
- RAVPN capacity assessment

b. System health-related insights

- High data plane CPU usage
- High snort CPU usage
- High data plane memory usage
- High snort memory usage

c. Anomaly detection-related insights

- Connections rate (CPS) anomaly
- Network throughput anomaly

Step 4 For each insight, expand the section and configure the available settings. Depending on the insight, you can select severity levels, configure thresholds, or specify forecast parameters.

Step 5 Click **Submit**.

After you enable the feature for the organization, you can view the detected anomalies in the **Summary** page, and the respective widgets are also displayed in the dashboard.

Best practices and recommendations insights

You can modify the preferences for **Best Practices & Recommendations**-related insights.

Note The setting for **Best Practices & Recommendations** is enabled by default.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Settings**.

Step 2 Click **Best Practices**.

Step 3 Enable the **Best Practices & Recommendations Analysis** toggle button to view the assessment categories, checks performed under each category, and the number of failed checks for each device.

Step 4 Enable the **Automatic assessment** toggle to activate automated report generation.

Step 5 Click **Submit**.

After you enable the feature for your tenant, you can view the detected anomalies in the **Summary** page, and the respective widget is displayed on the dashboard.

Policy change impact analyzer insights

You can modify the preferences for Policy Change Impact Analyzer-related insights.

Note The **Policy Change Impact Analyzer** toggle is enabled by default.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Settings**.

Step 2 Click **Policy Change Impact Analyzer** .

Step 3 Enable the **Policy Change Impact Analyzer** toggle.

This feature helps you review the potential impact of pending rule changes in Access Control policies before deployment. It currently supports Cloud-delivered Firewall Management Center-managed Firewall Threat Defense devices.

Step 4 Click **Submit**.

After the feature is enabled, you can access it by clicking **Insights & Reports** > **Policy Change Impact Analyzer**.

Policy change impact analyzer insights

You can modify the preferences for Policy Change Impact Analyzer-related insights.

Note The **Policy Change Impact Analyzer** toggle is enabled by default.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Settings**.

Step 2 Click **Policy Change Impact Analyzer** .

Step 3 Enable the **Policy Change Impact Analyzer** toggle.

This feature helps you review the potential impact of pending rule changes in Access Control policies before deployment. It currently supports Cloud-delivered Firewall Management Center-managed Firewall Threat Defense devices.

Step 4 Click **Submit**.

After the feature is enabled, you can access it by clicking **Insights & Reports** > **Policy Change Impact Analyzer**.

Operational insights

You can modify the preferences for operations-related insights.

Note The setting for **Software Update Planner** is enabled by default.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Settings**.

Step 2 Click **Operations**.

Step 3 Enable the **Software Update Planner** toggle. You can view the software versions running on your devices and our upgrade recommendations.

Step 4 Click **Submit**.

After you enable this feature for your tenant, detected insights are listed in the **Summary** page.

Application insights

You can modify the preferences for application outages- related insights.

Note The setting for **Application Insights** is enabled by default.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Settings**.

Step 2 Click **Application insights**.

Step 3 Enable the **Application outage insights** toggle button to receive insights and notifications when application outages occur.

Step 4 In the **Selected applications** list, search for an application and check the checkbox to include it for monitoring. Clear the checkbox to remove an application from monitoring.

Step 5 Click **Submit**.

After you enable the feature for your organization, you can view detected application outages on the **Application Insights** page, and corresponding insights are displayed on the **AgenticOps Summary** dashboard under the **Operations** category.

Frequently asked questions about AgenticOps

What is AgenticOps?

AgenticOps for firewalls leverages artificial intelligence (AI) and machine learning (ML) to streamline and enhance firewall management. By continuously analyzing the data, AgenticOps provides insights that help you:

- Optimize firewall policies and configurations.
- Detect misconfigurations before they impact performance.
- Improve feature adoption and adherence to best practices.
- Receive upgrade suggestions for Secure Firewall Threat Defense devices, including bugs fixes and security vulnerabilities fixes.

Are AgenticOps features available for all types of FMC-managed Firewall Threat Defense devices?

AgenticOps supports Firewall Threat Defense devices managed by Cloud-Delivered Firewall Management Center and On-Premises Firewall Management Centers registered with Cisco Security Cloud. Feature availability varies by management type. For details, see **AgenticOps Support for On-Premises Firewall Management Centers**.

Can onboarding AgenticOps fail?

If an onboarding failure occurs, open a support ticket with Cisco Technical Assistance Center (TAC).

Can AgenticOps Insights be disabled?

Opting out of AgenticOps stops data collection and disables insights. You can choose whether to retain or permanently delete historical data. For more information, see Disable AgenticOps insights, on page 349.

Can I disable specific AgenticOps modules?

Yes. From AgenticOps Settings, you can enable or disable modules such as Feature Adoption, Best Practices & Recommendations, and Operations. Disabled modules stop generating insights but do not affect device functionality.

Can I manually run Best Practices and Recommendations assessments?

Periodic assessments run automatically, but you can also run assessments manually at any time. For large-scale deployments with more than 50 devices, you must run assessments manually due to processing limits.

Does AgenticOps automatically make changes to my firewall configuration?

No. AgenticOps provides recommendations and guidance based on analysis of your environment. Configuration changes are not automatically applied unless you explicitly review and implement them.

Additional resources

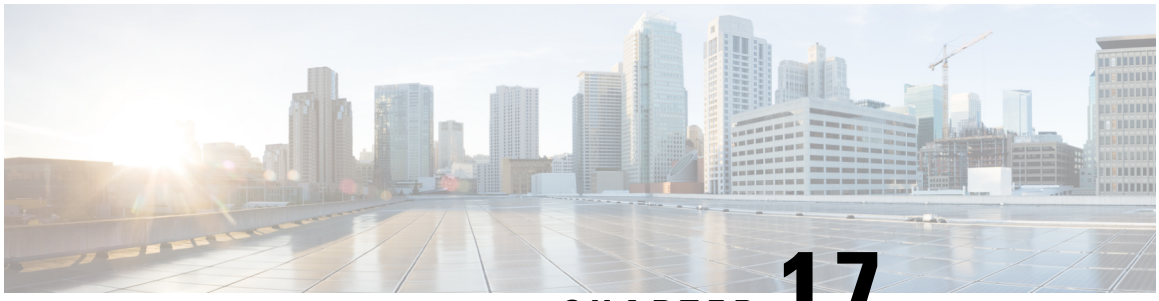
- From AIOps to AgenticOps: The Autonomous Evolution of Firewall Operations
- Managing Firewall complexity and Augmenting Effectiveness with AIOps for Cisco Firewall
- Security Cloud Control: Pioneering the Future of Security Management



PART V

Simplify Firewall Operations with Agent Workforce

- Introduction to Agent Workforce, on page 359



CHAPTER 17

Introduction to Agent Workforce

This chapter introduces Agent Workforce, an AI-driven operational framework in Security Cloud Control that helps you troubleshoot network issues, analyze operational behavior, and manage firewall policies using natural language interactions.

Agent Workforce uses specialized agents to assist with operational workflows such as VPN troubleshooting, traffic remediation, and policy analysis. You can interact with agents conversationally to investigate issues, review recommendations, and perform guided operational tasks.

- About Agent Workforce , on page 359
- Troubleshoot site-to-site VPN issues with VPN agent, on page 361
- Remediate traffic congestion with Elephant flow agent , on page 363
- Analyze firewall data with Firewall insight agent, on page 365

About Agent Workforce

Agent Workforce is an AI-driven operational framework that enables you to interact with specialized agents using natural language to troubleshoot issues, analyze network behavior, and manage firewall policies.

Instead of relying solely on manual investigation, command-line analysis, or complex configuration workflows, Agent Workforce allows you to describe operational problems or intended actions conversationally. Based on the user query, the appropriate agent analyzes telemetry, configurations, policies, and operational data to provide insights, recommendations, diagnostics, or policy actions. Each agent is optimized for a specific operational area and performs focused analysis based on its specialization.

Available agents in Agent Workforce

Agent Workforce includes specialized agents designed to assist with different operational and security workflows. Each agent focuses on a specific operational domain and uses relevant telemetry, protocol awareness, policy analysis, and contextual operational data to provide diagnostics, recommendations, and actionable outcomes.

When you enter a query, Agent Workforce interprets the intent and automatically routes the conversation to the most appropriate agent. You can describe operational issues or tasks in natural language without determining which agent should handle the request.

Note Currently, Agent Workforce capabilities are available only for deployments in the US, EU, and AUS regions.

Depending on the workflow, Agent Workforce may also present contextual quick actions to help you initiate common investigative or remediation tasks more efficiently. Currently, Agent Workforce includes the following agents:

- **VPN agent** helps troubleshoot site-to-site VPN connectivity, routing, tunnel negotiation, and secure traffic flow issues.
- **Elephant flow agent** helps analyze and remediate high-bandwidth traffic flows that may impact firewall performance and operational stability.
- **Firewall insight agent** helps analyze firewall inventory, configurations, policies, routing, and device status.

Manage Conversations in Agent Workforce

Use the **Conversations** page to create new conversations, monitor active investigations, and interact with specialized agents for troubleshooting, investigation, and remediation workflows.

Procedure

	Command or Action	Purpose
Step 1	In the left pane, click **Insights & Reports** > **Agent Workforce** > **Conversations.**	
Step 2	Review the list of available conversation threads.	Each conversation displays information such as the thread title, assigned agent, workflow status, creator, and creation date.
Step 3	To locate a specific conversation, use the **Search** field to search by keywords.	
Step 4	To narrow the conversation list, click **Filters** and filter conversations using one or more criteria.	
Step 5	To create a new conversation, click **New Conversation**.	
Step 6	Enter your query in natural language.	Agent Workforce automatically selects the most appropriate agent based on the operational intent of the request.
Step 7	Continue interacting with the assigned agent if additional investigation, troubleshooting, or remediation actions are required.	

Limitations and considerations

- Currently, Agent Workforce capabilities are available only for deployments in the US, EU, and AUS regions.
- AI-generated insights, recommendations, and remediation actions must always be reviewed and validated before deployment.
- The accuracy of generated responses depends on the availability, completeness, and correctness of telemetry, policies, device configurations, logs, objects, and operational context within the environment.
- Missing, incomplete, or unresolved configuration entities may prevent certain workflows, recommendations, remediation actions, or deployment operations from completing successfully.
- Agent responses and operational analysis may vary depending on the enabled platform capabilities, configured integrations, available data sources, and device reachability.
- Some workflows may require manual validation, administrative approval, or additional configuration updates before operational changes can be applied to managed devices or policies.
- AI-generated operational guidance is intended to assist administrators and must not replace standard security review, compliance validation, or change-management procedures.

Troubleshoot site-to-site VPN issues with VPN agent

The VPN agent is an intelligent troubleshooting agent that helps you diagnose and resolve site-to-site VPN connectivity and access issues. The agent analyzes operational telemetry, VPN configurations, routing behavior, packet flow, and access control policies to identify issues affecting secure connectivity between sites. It provides protocol-aware diagnostics across technologies such as BGP, EIGRP, OSPF, IPsec, and IKE to accelerate root-cause analysis and reduce manual troubleshooting effort.

Benefits of using the VPN agent

- Investigate VPN tunnel failures and connectivity disruptions.
- Analyze traffic flow between remote networks.
- Detect routing inconsistencies affecting VPN communication.
- Identify access control or policy-related packet drops.
- Perform packet-level troubleshooting and operational analysis.

How VPN agent works

When a query is submitted, the VPN agent correlates operational data across routing protocols, tunnel negotiation states, traffic flow behavior, and policy enforcement to identify the most likely root cause and provide actionable insights. Depending on the issue being investigated, the VPN agent can:

- Identify affected site-to-site VPN tunnels.
- Analyze IKE Phase 1 and Phase 2 negotiation states.

- Review tunnel establishment and peer connectivity status.
- Investigate routing behavior across the VPN tunnel.
- Evaluate traffic flow behavior between local and remote networks.
- Correlate logs, telemetry, and operational events to identify likely root causes.
- Generate operational findings, root-cause analysis, and remediation recommendations.



Note The VPN agent works only for Firewall Threat Defense site-to-site VPN configurations managed by Cloud-delivered Firewall Management Center or On-Premises Firewall Management Center.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Agent Workforce** > **Conversations**.

Step 2 Click **New Conversation**.

Step 3 Enter a query describing the VPN connectivity issue or traffic-flow problem.

Agent Workforce automatically assigns the request to the VPN agent based on the operational intent.

Step 4 Review the generated diagnostics, tunnel analysis, routing behavior, traffic flow analysis, and remediation recommendations.

Step 5 Continue refining the investigation by providing additional operational details such as affected sites, peer devices, tunnel names, or observed behavior.

Step 6 Review the recommended remediation actions and validate the operational impact before applying configuration changes.

After the analysis is completed, the VPN agent summarizes operational findings, identified root cause, impact assessment, additional observations, recommended remediation steps and next actions.

Best practices for using the VPN agent

To improve troubleshooting accuracy and reduce analysis time:

- Include source and destination networks whenever possible.
- Specify the affected VPN tunnel, site, or peer device if known.
- Clearly describe observed behavior such as packet drops, intermittent connectivity, or failed tunnel establishment.
- Refine prompts with additional operational context if further analysis is needed.
- Validate recommended remediation steps before deployment.

Example prompts and expected outcomes for the VPN agent

User intent	Example prompt	Expected outcome
Troubleshoot connectivity between remote sites	Troubleshoot why the connectivity is down between [remote site 1] and [remote site 2].	VPN agent identifies the relevant Site-to-Site VPN tunnel associated with the affected networks and performs guided troubleshooting analysis.
Troubleshoot a VPN tunnel outage	My VPN tunnel is down. Help me troubleshoot the issue.	VPN agent identifies available VPN tunnels, allows tunnel selection if needed, and analyzes tunnel state, negotiation behavior, routing, and connectivity conditions.
Investigate VPN traffic flow issues	Investigate why VPN traffic is not flowing across the [tunnel name].	VPN agent analyzes traffic forwarding behavior, routing conditions, and policy enforcement affecting traffic flow across the VPN tunnel.
Review configured VPN tunnels	How many VPN tunnels are configured and what state are they in?	VPN agent displays configured site-to-site VPN tunnels along with operational state, tunnel health, and connectivity status details.
Visualize VPN routing paths	Visualize the routing between various devices connected by my VPN tunnels.	VPN agent analyzes VPN topology, connected networks, and routing relationships, and provides communication paths across VPN tunnels.

Remediate traffic congestion with Elephant flow agent

Large traffic flows, commonly referred to as elephant flows, can consume excessive bandwidth, increase firewall processing overhead, and degrade application performance across the network. The Elephant flow agent is an intelligent remediation agent that helps you remediate high-bandwidth traffic flows that impact network performance and critical applications. It relies on a valid AgenticOps-generated insight and uses the device and application information provided in the insight to construct a trust rule and guide the remediation workflow. Unlike traditional troubleshooting workflows that require manual investigation across multiple dashboards and telemetry sources, the Elephant flow agent provides a guided remediation experience directly within the conversation workflow.



Note The Elephant flow agent remediates only high-bandwidth traffic flows identified in a valid AgenticOps-generated insight; it does not detect elephant flows.

Benefits of using the Elephant flow agent

- Analyze previously detected elephant flows affecting network performance.
- Investigate abnormal traffic spikes and sustained utilization.
- Identify applications or traffic flows contributing to congestion.
- Recommend remediation actions to reduce processing overhead and congestion.
- Stage policy updates associated with remediation workflows.

How Elephant flow remediation works

Before you begin

- Ensure that **Cloud-Delivered Firewall Management Center** is provisioned in your organization.
- Ensure that AgenticOps is onboarded in your organization.
- Ensure that **High traffic caused by elephant flow** is enabled under **Insights & Reports** > **AgenticOps Insights** > **Settings** > **Capacity & Traffic Analysis**. For more information, see Capacity and traffic insights.
- Ensure that valid AgenticOps-generated elephant flow insights exist. Remediation cannot begin without a valid insight.

Procedure

Step 1 In the left pane, click **Insights & Reports** >**Summary**.

Step 2 On the AgenticOps **Summary** page, select the **Traffic & Capacity** insights category.

Step 3 In the **High traffic caused by elephant flow** section, select the affected device.

Step 4 Review the detected elephant flow insight details, affected resources, applications, and traffic impact information.

Step 5 Click **Review proposed remediation** to open the remediation workflow in **Agent Workforce** > **Conversations**.

a. Alternatively, navigate to **Insights & Reports > Agent Workforce > Conversations**.

b. Click **New Conversation** and select **View elephant flow insights** to start the remediation workflow.

Note
Manual prompt entry is not required to start the workflow.

Step 6 Review the proposed remediation actions, operational warnings, shared policy impact, and success criteria.

Step 7 Approve or cancel the remediation workflow as required.

Step 8 Deploy the generated policy updates to the affected devices.

Note

Elephant flow events displayed in the Unified Event Viewer might not be available to the remediation agent if AgenticOps did not generate a corresponding insight.

Best practices for using the Elephant flow agent

To improve remediation accuracy and operational visibility:

- Review affected applications before approving remediation actions.
- Verify the operational impact of policy updates before deployment.
- Review shared policy impact warnings before approving remediation actions.
- Monitor application and network behavior after remediation is completed.

Analyze firewall data with Firewall insight agent

The Firewall insight agent helps you investigate the configuration and operational state of your managed firewalls. It provides information about device inventory, interfaces, routing, policies, licensing, compliance, software, and security configuration, and identifies conflicts, inconsistencies, and other issues that may require attention.

Benefits of using the Firewall insight agent

- View firewall inventory by device type, connectivity state, licensing status, compliance status, and onboarding state.
- Identify unlicensed or out-of-compliance devices.
- Identify access control or policy-related packet drops.
- Review routing configuration, BGP configurations, and route exchange between devices.
- Review interface status and device connectivity information.
- Review syslog and email alert configuration.
- Ask follow-up questions to investigate findings in greater detail.

How Firewall insight agent works

When you submit a query, Firewall insight agent identifies the relevant devices, configurations, policies, or operational data and returns the associated information. The agent can:

- Identify the devices, policies, routes, interfaces, or configurations related to the query.
- Correlate information across devices, policies, routing, security configuration, and operational state.
- Compare related configurations and identify conflicts, overlaps, redundancies, or inconsistencies.

- Explain when a device or configuration cannot be found and identify the information required to continue the investigation.
- Generate operational findings, root-cause analysis, and recommended next steps.

Procedure

Step 1 In the left pane, click **Insights & Reports** > **Agent Workforce** > **Conversations**.

Step 2 Click **New Conversation**.

Step 3 Enter a query about firewall inventory, device state, routing, policies, security configuration, or operational data.

Agent Workforce automatically assigns the request to an agent based on the operational intent.

Step 4 Review the returned firewall information, analysis, and key observations.

The Firewall insight agent analyzes firewall data and reports findings, comparisons, and recommended next steps. It does not modify configurations, create or deploy policies, or approve changes.

Step 5 Continue the investigation by providing additional context, such as device names, policy names, object names, interfaces, or other relevant identifiers.

Step 6 Review the recommended next steps and validate the operational impact before applying configuration changes.

After the analysis is completed, Firewall insight agent summarizes the operational findings, identified root cause, impact, additional observations, and recommended next steps.

Best practices for using the Firewall insight agent

To improve troubleshooting accuracy and reduce analysis time:

- Include device names, policy names, object names, or other identifiers when available.
- Clearly describe observed behavior such as packet drops, intermittent connectivity, unexpected policy behavior, or configuration inconsistencies.
- Provide additional operational context if further analysis is needed.
- Review recommended next steps and validate their operational impact before deployment.

Example prompts and expected outcomes for the Firewall insight agent

User intent	Example prompt	Expected outcome
Review firewall inventory	What devices do I have in my firewall inventory?	The agent summarizes devices by type and connectivity state. It can also provide software versions, addresses, manager associations, licensing status, and compliance status.

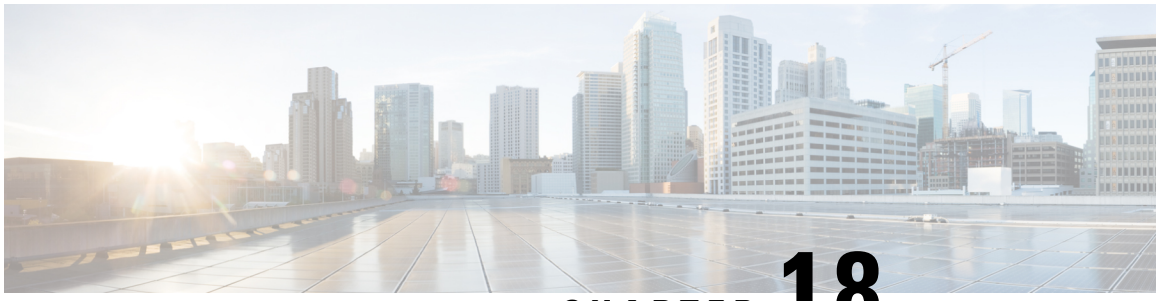
User intent	Example prompt	Expected outcome
Review interfaces	List all interfaces and their status for `[device name]`.	The agent lists the device interfaces and their status. If the device cannot be found, it explains the lookup failure.
Review BGP configuration	Show the current BGP configuration on `[device name]`.	The agent returns BGP neighbors, local and remote AS numbers, peering type, tunnel interfaces, redistribution, and related configuration details.
Review licensing and compliance	Which Firewall Threat Defense devices are unlicensed or out of compliance?	The agent lists affected devices and provides their license and compliance status, software version, and other relevant device details.
Review software updates	Which Firewall Threat Defense devices have a recommended software update?	The agent identifies devices with available updates and provides current versions, target versions, fixes, and affected defects or threats.



PART VI

Health and Monitoring

- Health, on page 371
- Troubleshooting, on page 419



CHAPTER 18

Health

The following topics describe how to use health monitoring:

- Requirements and prerequisites for health monitoring, on page 371
- Health monitoring in Cloud-Delivered Firewall Management Center, on page 371
- Health policies, on page 383
- Device exclusion in health monitoring, on page 389
- Health monitor alerts, on page 393
- About the health monitor, on page 398
- Health event views, on page 407
- Sending device health metrics to OpenConfig, on page 410
- system auditing, on page 416

Requirements and prerequisites for health monitoring

This reference provides the requirements and prerequisites needed to implement health monitoring functionality.

Model support

Any

Supported domains

Any

Health monitoring in Cloud-Delivered Firewall Management Center

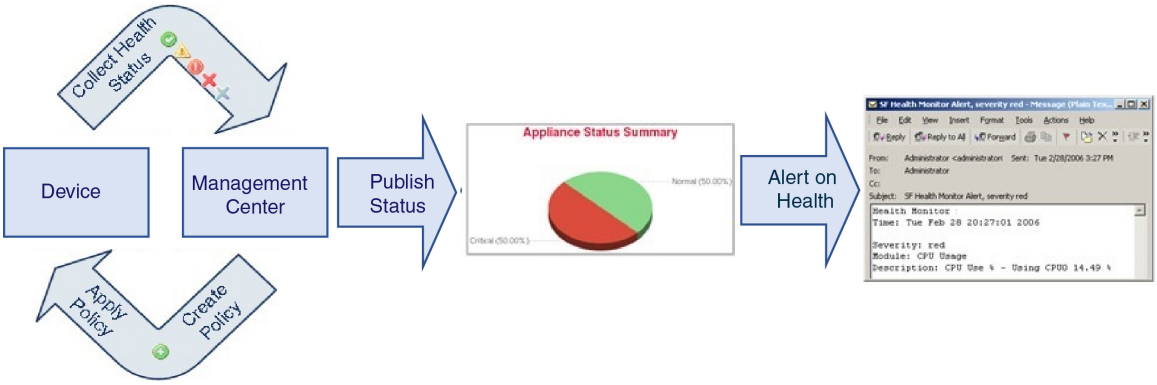
Health monitoring is a system capability that

- tracks various health indicators to ensure hardware and software are working correctly,
- checks the status of critical functionality across your deployment, and
- enables configuration of test frequency for alerting and time series data collection.

Health monitoring capabilities

Health monitoring includes two key capabilities: health status monitoring through health policies and health modules, and time series data collection for metrics analysis.

With time series data collection, you can configure how frequently metric data is collected for the device and its health modules. The device monitor reports these metrics in several predefined health monitor dashboards by default. This metric data is collected for analysis. Some alerts are generated from metric values when they cross the thresholds specified in the health policy.



You can use the health monitor to create a collection of tests, referred to as a *health policy*, and apply the health policy to one or more appliances. The tests, referred to as *health modules*, are scripts that test for the criteria you specify. You can modify a health policy by enabling or disabling tests or by changing test settings, and you can delete health policies that you no longer need. You can also suppress messages from selected appliances by excluding them.

The health monitoring system run the tests in a health policy at the configured intervals. You can also run all tests, or a specific test, on demand. The health monitor collects health events based on the test conditions configured.

Note All appliances automatically report their hardware status via the Hardware Alarms health module. The Cloud-Delivered Firewall Management Center also automatically reports status using the modules configured in the default health policy. Some health modules, such as the Appliance Heartbeat module, run on the Cloud-Delivered Firewall Management Center and report the status of the Cloud-Delivered Firewall Management Center's managed devices. For the health modules to provide managed device status, you must deploy all health policies to the device.

You can use the health monitor to access health status information for the entire system and for a particular appliance. In a multi-domain deployment, you can view the health status summary for a device in the domain where the device is located.

A hexagon-shaped widget and status tables on the **Health Status** page provide a visual summary of the status of all appliances on your network, including the Cloud-Delivered Firewall Management Center. Individual appliance health monitors let you drill down into health details for a specific appliance.

Fully customizable event views allow you to quickly and easily analyze the health status events gathered by the health monitor. These event views allow you to search and view event data and to access other information that may be related to the events you are investigating. For example, if you want to see all the occurrences of CPU usage with a certain percentage, you can search for the CPU usage module and enter the percentage value.

You can also configure alerting in response to health events. A *health alert* is an association between a standard alert and a health status level. SNMP and syslog are not supported as the primary methods for sending health alerts to external systems. Use email or webhook alert responses for external health alerting. For example, if you want to make sure an appliance never fails due to hardware overload, you can set up an email alert. You can then create a health alert that triggers an email alert whenever CPU, disk, or memory usage reaches the Warning level you configure in the health policy applied to that appliance. You can set alerting thresholds to minimize the number of repeating alerts you receive.

Note The health monitoring can take 5–6 minutes from the occurrence of the health event to generate the health alert.

You can also generate troubleshooting files for an appliance if you are asked to do so by Support.

High availability pair

In a Cloud-Delivered Firewall Management Center high availability deployment, the active peer creates a health monitor page that uses REST APIs to show detailed metric-based information. The standby peer creates the health monitor page that shows the alert information and provide a visual summary of the status of all appliances on your network using pie charts and status tables. The standby peer does not display the metric-based information.

Health modules

Describes health modules that monitor system conditions and provide alerts or metrics for managing device health and performance.

Health modules, or *health tests,* test for the conditions that you specify in a health policy.

Health monitoring includes both health alerts or tests and metric or time series data collection. The two types of health modules are alert modules and metrics modules. Alerts modules (sometimes called *legacy* modules) monitor system infrastructure and report only the health status. When the conditions specified in the health policy for these monitored systems are met, these modules raise health alerts. Metrics modules (sometimes called *telegraf* modules) collect statistics (sometimes called *time series data*) that you can view on the health monitoring dashboard. Some alerts are also generated from metric values when they cross the thresholds specified in the health policy. You can create custom dashboards with your preferred health metrics, allowing you to monitor statistics or troubleshoot appliance health issues.

Note The health alerts generated from the Secure Firewall 200 series device is limited to the essential health modules, to optimize performance and ensure effective resource utilization. For more information about the available health modules, refer to Health alerts for Secure Firewall 200 Series device, on page 394.

Table 19: Device Health Modules

Module	Type	Description
AMP Connection Status	Metrics	The module alerts if the device cannot connect to the AMP cloud or Cisco AMP Private Cloud after an initial successful connection, or if the private cloud cannot contact the public AMP cloud. Disabled by default.

Module	Type	Description
AMP Threat Grid Connectivity	Metrics	The module alerts if the device cannot connect to the AMP Threat Grid cloud after an initial successful connection.
ASP Drop	Alert	Monitors the connections dropped by the data plane accelerated security path.
Automatic Application Bypass	Alert	Monitors bypassed detection applications.
Certificate Monitoring	Alert	Alerts when service authentication certificates are near expiration or have expired, based on a configurable threshold (in days). This alert helps you to identify certificates that are about to expire and renew them before a service disruption occurs.
Chassis Environment Status	Alert	Monitors chassis parameters such as fan speed and chassis temperature, and enables you to set a warning threshold and critical threshold for temperature. The **Critical Chassis Temperature (Celsius)** default value is `85`. The **Warning Chassis Temperature (Celsius)** default value is `75`.
Cluster/HA Failover Status	Alert	For threat defense clusters, alerts when a unit joins, leaves, or is elected primary.
Configuration Resource Utilization	Alert	Alerts if the size of your deployed configurations puts a device at risk of running out of memory. The alert shows you how much memory your configurations require, and by how much this exceeds the available memory. If this happens, reevaluate your configurations. You may be able to reduce the number or complexity of access control rules or intrusion policies.
Connection Statistics	Metrics	Monitors connection statistics and NAT translation counts. For standby Firewall Threat Defense devices in a high availability pair, this widget reflects only the statistics of connections replicated from the active device.
CPU Usage (per core)	Metrics	Alerts when CPU core use exceeds a configurable threshold.
Critical Process Statistics	Metrics	Monitors the state of critical processes, their resource consumption, and the restart counts.
CPU Usage Date Plane	Metrics	Alerts when data plane CPU use exceeds a configurable threshold.
Memory Usage Data Plane	Metrics	Alerts when data plane memory use exceeds a configurable threshold.
Deployed Configuration Statistics	Metrics	Monitors statistics about the deployed configuration, such as the number of ACEs and IPS rules.
Disk Status	Alert	Alerts if there is an issue with the hard disk or RAID controller. If this module alerts, contact Cisco TAC. This will prevent upgrade.

Module	Type	Description
Disk Usage	Metrics	This module compares disk usage on the appliance's hard drive to the limits configured for the module and alerts when usage exceeds the thresholds configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds. Use the Disk Usage health status module to monitor disk usage for the `/` and `/volume` partitions on the appliance and track draining frequency. Although the disk usage module lists the `/boot` partition as a monitored partition, the size of the partition is static so the module does not alert on the boot partition. Use the **Clear disk space** option to free up disk space by removing the temporary files from your threat defense device. For more information, refer to Clear disk space, on page 432
File System Integrity Check	Alert	This module performs a file system integrity check and runs if the system has CC mode or UCAPL mode enabled, or if the system runs an image signed with a DEV key.
Firewall Threat Defense HA	Alert	Alerts if a threat defense high availability pair is split brain.
Firewall Threat Defense Platform Faults	Alert	Monitors Secure Firewall 1000 /3100 /4200 platform faults and generate health alerts for the faults. A platform fault represents a failure in the Firewall Threat Defense instance or an alarm threshold that has been raised. During the lifecycle of a platform fault, it can change from one state or severity to another. Each fault includes information about the operational state of the affected object at the time the fault was raised. If the fault is transitional and the failure is resolved, then the object transitions to a functional state. For more information, refer to *Cisco Firepower 1000/2100 FXOS Faults and Error Messages Guide*.
Flow Offload Statistics	Metrics	Monitors hardware flow offload.
Hardware Alarms	Alert	This module determines if hardware needs to be replaced on a physical managed device and alerts based on the hardware status. It also reports on the status of hardware-related daemons.
Identity Limits Monitor (supported only on Secure Firewall 200 Series devices)	Alert	Alerts when the device identity-related mappings and user-to-group mappings exceed the normal limit. Device identity-related mappings include user sessions, SGT Exchange Protocol (SXP) mappings, and dynamic object mappings. For more information, refer to Identity-limits-monitor.
Inline Link Mismatch Alarms	Alert	Alerts if inline pair interfaces negotiate different speeds.

Module	Type	Description
Interface Status	Alert	Determines if the device currently collects traffic and alerts based on the traffic status of physical interfaces and aggregate interfaces. For physical interfaces, the information includes interface name, link state, and bandwidth. For aggregate interfaces, the information includes interface name, number of active links, and total aggregate bandwidth. **Note** • This module also monitors the high availability standby device traffic flow. Though it is known that the standby device would not be receiving any traffic yet, the Cloud-Delivered Firewall Management Center alerts that the interface is not receiving any traffic. The same alerting principle is applied when traffic is not received by some of the subinterfaces on a port channel. • This module displays the traffic rates according to the values from Lina. However, if you use the **show interface** CLI command to know the interface statistics of your device, the input and output rates in the CLI command result can be different from the traffic rates that appear in the **Interface** widget. The sampling intervals of Lina and the Cloud-Delivered Firewall Management Center interface statistics are different. Due to the difference in sampling interval, throughput values in the Cloud-Delivered Firewall Management Center GUI can be different from the throughput values appears in the device CLI result. • Note that traffic rates in the **Interface Traffic Rate** widget (**Overview** > **Dashboards heading** > **Dashboard** page) can be different as it displays the input and output rates from Snort.
Intrusion and File Event Rate	Alert	Alerts if intrusion events per second exceed a configurable threshold. We recommend a warning threshold of 1.5 times your average intrusion event rate, and a critical threshold of 2.5 times. For example, for an average event rate on network segment of 20 events per second, we recommend a warning value of 30 and a critical value of 50. The critical limit must be lower than 1000, and higher than the warning limit. Event rates for your devices are available on **System** (⚙) > **Monitoring** > **Statistics**. If the rate is zero, the Snort process may be down or the device may not be sending events.
Link State Propagation	Alert	For the ISA 3000, alerts when an interface in a inline set fails.
Memory Usage	Alert	Alerts when memory use exceeds configurable thresholds. For appliances with more than 4 GB of memory, the preset alert thresholds are based on a formula that accounts for proportions of available memory likely to cause system problems. On >4 GB appliances, because the interval between Warning and Critical thresholds may be very narrow, it is recommended that you manually set the **Warning Threshold %** value to `50`. This will further ensure that you receive memory alerts for your appliance in time to address the issue. Complex access control policies and rules can command significant resources and negatively affect performance.

Module	Type	Description
Network Card Reset	Alert	Alerts when a network card restarts due to hardware failure.
NTP Statistics	Metrics	Monitors NTP synchronization status. Disabled by default.
Firewall Management Center Access Configuration Changes	Alert	Monitors configuration changes made on the Cloud-Delivered Firewall Management Center directly using the **configure network management-data-interface** command. This module alerts when there is a conflict between the existing Cloud-Delivered Firewall Management Center configuration and the out of band configuration changes made.
Process Status	Alert	Alerts when processes on the appliance exit or terminate outside of the process manager. If a process is deliberately exited outside of the process manager, the module status changes to Warning and the health event message indicates which process exited, until the module runs again and the process has restarted. If a process terminates abnormally or crashes outside of the process manager, the module status changes to Critical and the health event message indicates the terminated process, until the module runs again and the process has restarted.
Routing Statistics	Metrics	Monitors the current state of routing table.
Snort 3 Statistics	Metrics	Collects Snort 3 statistics for events, flows, and packets. This module also monitors metrics for sending advanced logging events and generates the following alerts: • **Advanced logging events to syslog servers were dropped**: This alert appears when syslog messages are dropped due to memory overflow. • **Advanced Logging events failed to transmit to syslog servers**: This alert appears when the syslog messages failed to transmit due to a connection issue with the syslog server or a configuration error. Check your syslog server status and syslog configuration in Cloud-Delivered Firewall Management Center.
CPU Usage Snort	Metrics	This module checks that the average CPU usage of the Snort processes on the device is not overloaded and alerts when CPU usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`.
Snort Identity Memory Usage	Alert	Enables you to set a warning threshold for Snort identity processing and alerts when memory usage exceeds the level configured for the module. The **Critical Threshold %** default value is `80`. This health module specifically keeps track of the total space used for the user identity information in Snort. It displays the current memory usage details, the total number of user-to-IP bindings, and user-group mapping details. Snort records these details in a file. If the memory usage file is not available, the Health Alert for this module displays *Waiting for data*. This could happen during a Snort restart due to a new install or a major update, switch from Snort 2 to Snort 3 or back, or major policy deployment. Depending on the health monitoring cycle, and when the file is available, the warning disappears, and the health monitor displays the details for this module with its status turned Green.

Module	Type	Description
Memory Usage Snort	Metrics	This module checks the percentage of allocated memory used by the Snort process and alerts when memory usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`.
Snort Reconfiguring Detection	Metrics	Alerts if a device reconfiguration has failed. This module detects reconfiguration failure for both Snort 2 and Snort 3 instances.
Snort Statistics	Metrics	Monitors Snort statistics for events, flows, and packets.
SSE Connection Status	Metrics	The module alerts if the device cannot connect to the security services exchange cloud after an initial successful connection. Disabled by default.
CPU Usage System	Metrics	This module checks that the average CPU usage of all system processes on the device is not overloaded and alerts when CPU usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`.

Module	Type	Description
Threat Data Updates on Devices	Alert	Certain intelligence data and configurations that devices use to detect threats are updated on the Cloud-Delivered Firewall Management Center from the cloud every 30 minutes. This module alerts you if this information has not been updated on the devices within the time period you have specified. Note that the Secure Firewall 200 series device does not maintain a local URL database and supports Cloud Only lookup. Local URL database related alerts are not available for this device type. Monitored updates include: • Local URL category and reputation data • Security Intelligence URL lists and feeds, including global Block and Do Not Block lists and URLs from Threat Intelligence Director • Security Intelligence network lists and feeds (IP addresses), including global Block and Do Not Block lists and IP addresses from Threat Intelligence Director • Security Intelligence DNS lists and feeds, including global Block and Do Not Block lists and domains from Threat Intelligence Director • Local malware analysis signatures (from ClamAV) • SHA lists from Threat Intelligence Director, as listed on the **Objects > Object Management > Security Intelligence > Network Lists and Feeds** page • Dynamic analysis settings configured on the **Administration** > **Dynamic Attributes Connector** page • Threat Configuration settings related to expiration of cached URLs, including the Cached URLs Expire setting on the **Integration** > **Other Integrations** > **Cloud Services** page. (Updates to the URL cache are not monitored by this module.) • Communication issues with the Cisco cloud for sending events. Refer to **Cisco Cloud** box on the **Integration** > **Other Integrations > Cloud Services** page. **Note** Threat Intelligence Director updates are included only if TID is configured on your system and you have feeds. By default, this module sends a warning after 1 hour and a critical alert after 24 hours. If this module indicates failure on the Cloud-Delivered Firewall Management Center or on any devices, verify that the Cloud-Delivered Firewall Management Center can reach the devices.
VPN Statistics	Metrics	Monitors site-to-site and remote access VPN tunnels between Firewall Threat Defense devices.
XTLS Counters	Metrics	Monitors XTLS/SSL flows, memory and cache effectiveness. Disabled by default.

Table 20: Management Center Health Modules

Module	Type	Description
Secure Endpoint Status	Alert	The module alerts if the Cloud-Delivered Firewall Management Center cannot connect to the AMP cloud or Cisco AMP Private Cloud after an initial successful connection, or if the private cloud cannot contact the public AMP cloud. It also alerts if you deregister an AMP cloud connection using the Secure Endpoint management console.
AMP for Firepower Status	Alert	Alerts if: • The Cloud-Delivered Firewall Management Center cannot contact the AMP cloud (public or private) or the Secure Malware Analytics Cloud or Appliance, or the AMP private cloud cannot contact the public AMP cloud. • The encryption keys used for the connection are invalid. • A device cannot contact the Secure Malware Analytics Cloud or Secure Malware Analytics Appliance to submit files for dynamic analysis. • An excessive number of files are detected in network traffic based on the file policy configuration. If your Cloud-Delivered Firewall Management Center loses connectivity to the Internet, the system may take up to 30 minutes to generate a health alert.
Appliance Heartbeat	Alert	This module determines if an appliance heartbeat is being heard from the appliance and alerts based on the appliance heartbeat status.
Dynamic Attributes Connector Status	Alert	This module monitors the operational status of the Dynamic Attributes Connector (CSDAC) service in Cloud-Delivered Firewall Management Center and raises an alert if the CSDAC service is not available. If the CSDAC service is not being used by a related service, try disabling and re-enabling it in the Cloud-Delivered Firewall Management Center UI, and check the health alert again after some time. For more information on how to enable CSDAC service, refer to Enable the dynamic attributes connector. If a related service such as an Azure AD realm is configured, you cannot disable the CSDAC service directly. First delete the Azure AD realm and then disable and re-enable the CSDAC service. For more information, refer to Configure realms and directories.
Firewall Threat Defense HA Status	Alert	Alerts if a threat defense high availability pair is split brain.
Hardware Statistics	Metrics	Monitors Cloud-Delivered Firewall Management Center hardware: fan speed, temperature, and power supply. Alerts when values exceed configurable thresholds.
Health Monitor Process	Alert	Monitors the health process itself, and alerts if there have been no health events in some number of minutes (configurable).
ISE Connection Monitor	Alert	This module monitors the status of the server connections between the Cisco Identity Services Engine (ISE) and the Cloud-Delivered Firewall Management Center. ISE provides additional user data, device type data, device location data, SGTs (Security Group Tags), and SXP (Security Exchange Protocol) services.

Module	Type	Description
Local Malware Analysis	Alert	This module monitors ClamAV updates for Local Malware Analysis.
Passive Identity Agent Status Monitor	Alert	Displays connection errors between the Cloud-Delivered Firewall Management Center and the machine on which it's installed. The passive identity agent periodically sends updates to the Cloud-Delivered Firewall Management Center. This health alert is displayed if the Cloud-Delivered Firewall Management Center receives an error from the passive identity agent or if the Cloud-Delivered Firewall Management Center has not received an update or response for five minutes or longer. Try verifying the connection between the two, restarting the passive identity agent software, and checking the health alert again in a few minutes. If issues persist, check the configuration at **Integration** > **Other Integrations** > **Identity Sources**.
Process Status	Alert	Alerts when processes on the appliance exit or terminate outside of the process manager. If a process is deliberately exited outside of the process manager, the module status changes to Warning and the health event message indicates which process exited, until the module runs again and the process has restarted. If a process terminates abnormally or crashes outside of the process manager, the module status changes to Critical and the health event message indicates the terminated process, until the module runs again and the process has restarted.
Security Intelligence	Alert	Alerts if Security Intelligence is in use and the Cloud-Delivered Firewall Management Center cannot update a feed, or feed data is corrupt or contains no recognizable IP addresses. Refer to the Threat Data Updates on Devices module.
Smart License Monitor	Alert	Monitors Smart Licensing status and alerts if: • There is a communication error between the Smart Licensing Agent (Smart Agent) and the Smart Software Manager. • The Product Instance Registration Token has expired. • The Smart License usage is out of compliance. • The Smart License authorization or evaluation mode has expired.
Talos Connectivity Status	Alert	Monitors connectivity with Talos, required to download URL filtering and event enrichment data.

Module	Type	Description
Threat Data Updates on Devices	Alert	Certain intelligence data and configurations that devices use to detect threats are updated on the Cloud-Delivered Firewall Management Center from the cloud every 30 minutes. This module alerts you if this information has not been updated on the devices within the time period you have specified. **Note** The Secure Firewall 200 does not maintain a local URL database and supports only cloud-based lookups. Local URL database-related alerts are not available for this device type. Monitored updates include: • Local URL category and reputation data. • Security Intelligence URL lists and feeds, including global Block and Do Not Block lists and URLs from Threat Intelligence Director. • Security Intelligence network lists and feeds (IP addresses), including global Block and Do Not Block lists and IP addresses from Threat Intelligence Director. • Security Intelligence DNS lists and feeds, including global Block and Do Not Block lists and domains from Threat Intelligence Director. • Local malware analysis signatures (from ClamAV). • SHA lists from Threat Intelligence Director, as listed on the **Objects** > **Object Management** > **Security Intelligence** > **Network Lists and Feeds** page. • Dynamic analysis settings configured on the **Administration** > **Dynamic Attributes Connector** page. • Threat Configuration settings related to expiration of cached URLs, including the Cached URLs Expire setting on the **Integration** > **Other Integrations** > **Cloud Services** page. (Updates to the URL cache are not monitored by this module.) • Communication issues with the Cisco cloud for sending events. Refer to the **Cisco Cloud** box on the **Integration** > **Other Integrations > Cloud Services** page. **Note** Threat Intelligence Director updates are included only if TID is configured on your system and you have feeds. By default, this module sends a warning after 1 hour and a critical alert after 24 hours. If this module indicates failure on the Cloud-Delivered Firewall Management Center or on any devices, verify that the Cloud-Delivered Firewall Management Center can reach the devices.
Time Synchronization Status	Alert	This module tracks the synchronization of a device clock that obtains time using NTP with the clock on the NTP server and alerts if the difference in the clocks is more than ten seconds.

Module	Type	Description
Unresolved Groups Monitor	Alert	Monitors Foreign Security Principals (FSPs) that are groups used in policies. Security principals are Active Directory objects, like authenticated user groups, to which security can be applied in access control policies. This module generates a warning alert for unresolved groups that exist but are not used in policies, and a critical alert for unresolved groups that are used in policies.
URL Filtering Monitor	Alert	Monitors connectivity with the Cisco cloud, which is required for downloading URL filtering data and doing URL filtering lookups.
Zero-Touch Provisioning	Alert	Alerts if there is a failure when registering a device using the serial number. It also shows errors related to zero-touch provisioning capable Cloud-Delivered Firewall Management Center s in high availability.

Configure health monitoring

Health monitoring enables proactive management of network appliances by tracking their operational status using customizable policies and alert mechanisms. You can configure specific monitoring policies for different appliance types, enabling only the relevant tests for each type.

Procedure

Step 1 Determine which health modules you want to monitor as discussed in Health modules, on page 373.

You can set up specific policies for each kind of appliance, enabling only the appropriate tests for that appliance.

Tip
To quickly enable health monitoring without customizing the monitoring behavior, you can apply the default policy provided for that purpose.

Step 2 Apply a health policy to each appliance where you want to track health status as discussed in Create a health policy, on page 385.

Step 3 (Optional.) Configure health monitor alerts as discussed in Create a health monitor alert, on page 396.

You can set up email or webhook alerts that trigger when the health status level reaches a particular severity level for specific health modules.

Health policies

A health policy is a configuration framework that

- enables you to define test criteria for monitoring the health of appliances,
- allows selective enabling of health modules and critical alert settings, and

- provides flexibility to apply, customize, or use default policies for each appliance.

Health policy configuration options

When you configure a health policy, you decide whether to enable each health module for that policy. You also select the criteria that control which health status each enabled module reports each time it assesses the health of a process.

You can enable or disable health alerts for individual attributes within a health module. By adjusting the settings for health alerts at the attribute level, you can reduce the number of health alerts and streamline your focus on the most critical health alerts without interrupting data collection. You can continue to monitor the attribute from your **Health Monitor** dashboard. When you enable a health module in the health policy, the health alert configuration for all the attributes within that health module gets enabled by default.

You can create one health policy that can be applied to every appliance in your system, customize each health policy to the specific appliance where you plan to apply it, or use the default health policy provided for you.

Note When you register an appliance, the Cloud-Delivered Firewall Management Center automatically assigns it the default health policy. To disassociate a health policy from an appliance, you must first associate a different health policy with it. An appliance must have at least one health policy assigned.

Default health policy

The Cloud-Delivered Firewall Management Center setup process creates and applies an initial health policy, in which most—but not all—available health modules are enabled. The system also applies this initial policy to devices added to the Cloud-Delivered Firewall Management Center.

This initial health policy is based on a default health policy. You cannot view or edit the default health policy, but you can copy it when creating a custom health policy.

You can create a custom health policy and set it as your default health policy. The Cloud-Delivered Firewall Management Center applies the default health policy to any managed device when you add the device to the Cloud-Delivered Firewall Management Center. Note that you cannot delete the health policy that you have set as the default. For detailed instructions for setting a default health policy, refer to Set a default health policy, on page 388.

Upgrades and the default health policy

When you upgrade the Cloud-Delivered Firewall Management Center, any new health modules are added to all health policies, including the initial health policy, default health policy, and any other custom health policies. Usually, new health modules are added in an enabled state.

Note For a new health module to begin monitoring and alerting, reapply health policies after upgrade.

Create a health policy

If you want to customize a health policy to use with your appliances, you can create a new policy. The settings in the policy initially populate with the settings from the health policy you choose as a basis for the new policy. You can edit the policy to specify your preferences, such as enable or disable modules within the policy, change the alerting criteria for each module as needed, and specify the run time intervals.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click **Create Policy**.

Step 3 Enter a name for the policy.

Note that the following names are reserved for the default policies, and you cannot create a health policy using these names:

- Default Device Policy
- Default Health Policy

Step 4 Choose the existing policy that you want to use as the basis for the new policy from the **Base Policy** drop-down list.

Step 5 Enter a description for the policy.

Step 6 Choose **Save**.

What to do next

- Apply the health policy on devices as described in Apply a health policy, on page 385.
- Edit the policy to specify the module-level policy settings as described in Edit a health policy, on page 386.

Apply a health policy

When you apply a health policy to an appliance, the health tests for all the modules you enabled in the policy automatically monitor the health of the processes and hardware on the appliance. Health tests then continue to run at the intervals you configured in the policy, collecting health data for the appliance and forwarding that data to the Cloud-Delivered Firewall Management Center.

If you enable a module in a health policy and then apply the policy to an appliance that does not require that health test, the health monitor reports the status for that health module as disabled.

To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored.

When you apply a different policy to an appliance that already has a policy applied, expect some latency in the display of new data based on the newly applied tests.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click the **Deploy health policy** (📥) next to the policy you want to apply.

Step 3 Choose the appliances where you want to apply the health policy.

Note
An appliance must have at least one health policy assigned to it. To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored. To disassociate a health policy from an appliance, you must first associate a different health policy with it.

Step 4 Click **Apply** to apply the policy to the appliances you chose.

The health policy is successfully applied to the selected appliances, and monitoring begins automatically.

What to do next

- Optionally, monitor the task status; see View task messages, on page 426.

  Monitoring of the appliance starts when the policy is successfully applied.

Edit a health policy

Edit a health policy when you need to adjust which modules are monitored, change critical or warning thresholds, or reassign policies to different devices.

Before you begin

- Ensure you have access to Cloud-Delivered Firewall Management Center.
- Identify the health policy you need to modify.

Follow these steps to edit a health policy:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click **Edit** (✎) next to the policy you want to modify.

Step 3 To edit the policy name and its description, click the **Edit** (✎) icon provided against the policy name.

Step 4 The **Health Modules** tab displays all the device modules and its attributes. Configure your health modules using the following actions:

- Click the toggle button that is provided against the module and its attributes—turn on () or turn off () to enable or disable testing of health status respectively.
- To execute a bulk enable or disable testing on the health modules, click the **Select All** toggle button
- After you enable a health module and when available, use the check box next to an individual attribute within a health module to enable or disable health alerts for that attribute. Note that disabling health alerts for an attribute does not stop the collection of metrics for that attribute. You can continue to monitor these attributes from the **Health Monitor** dashboard.

Note

- The modules and attributes are flagged with the supporting appliances—Firewall Threat Defense, Cloud-Delivered Firewall Management Center, or both.
- You cannot choose to include or exclude the individual attributes of CPU and Memory modules.

For information on the modules, refer to Health modules, on page 373.

Step 5 Where appropriate, set the **Critical** and **Warning** threshold percentages.

Step 6 In the **Settings** tab, enter the relevant values in the fields:

- **Health Module Run Time Interval**: The frequency for running the health modules. The minimum interval is 5 minutes.
- **Metric Collection Interval**: The frequency of collecting the time series data on the device and its health modules. The device monitor reports these metrics in several predefined health monitor dashboards by default. For detailed information on the dashboard, refer to View the device health monitor, on page 401. The metric data is collected for analysis and hence no alerting is associated with it.
- **OpenConfig Streaming Telemetry**—Configure a health metrics telemetry stream from the Firewall Threat Defense devices to an external data collection system which uses the vendor neutral, OpenConfig model. For more information, refer to Configure OpenConfig Streaming Telemetry.

Step 7 To view and modify the devices to which the policy is assigned, do the following:

a) Click **Policy Assignments & Deploy**.
b) From the **Available Devices** list, click the + icon next to the device to which you want to assign the health policy.
c) Click **Apply**.

Note
The UI prompts you to save unsaved changes before continuing with assignment and deployment.

Alternatively, you can apply the health policy to your appliance as described in Apply a health policy, on page 385

Apply the health policy to each appliance where you want to track health status. When you apply the health policy to an appliance, all the modules you enabled in the policy monitor the health of the processes and hardware on the appliance, and forwards that data to the Cloud-Delivered Firewall Management Center.

Step 8 Click **Save**.

Set a default health policy

You can set a user-created health policy as your default health policy. The Cloud-Delivered Firewall Management Center applies the default health policy to any managed device when you add the device to the Cloud-Delivered Firewall Management Center.

Note Setting a new default health policy does not affect the health policy assigned to devices that are already registered.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click the **More Actions** (⋮) icon next to the health policy that you want to set as the default, and click **Set as Default**.

Step 3 Click **Proceed**.

Delete a health policy

You can delete health policies that you no longer need. However, an appliance must have at least one health policy assigned to it. If you delete a policy that is still applied to an appliance, the policy settings remain in effect until you apply a different policy. In addition, if you delete a health policy that is applied to a device, any health monitoring alerts in effect for the device remain active until you disable the underlying associated alert response.

Note To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click **Delete** (🗑) next to the policy that you want to delete.

Step 3 Click **Delete health policy** to delete it.

Note that you cannot delete a health policy that is set as the default policy. Set another health policy as the default before attempting to delete the current default policy. For more information, refer to Set a default health policy, on page 388.

Device exclusion in health monitoring

A device exclusion is a health monitoring configuration that

- prevents deliberate device or module outages from affecting the overall health summary,
- allows administrators to temporarily disable health monitoring for selected devices, Cloud-Delivered Firewall Management Center where supported, or specific health modules, and
- ensures health events generated during exclusion are marked as disabled and do not affect health calculations.

Device exclusion behavior

In the course of normal network maintenance, you may disable managed devices or make them temporarily unavailable. In Cloud-Delivered Firewall Management Center, you can also exclude Cloud-Delivered Firewall Management Center health monitoring where applicable. Because those outages or expected conditions are deliberate, you may not want their health status to affect the overall health summary.

When you disable health monitoring status, health events are still generated, but they have a disabled status and do not affect the health status for the health monitor. If you remove the device or module from the excluded list, the events that were generated during the exclusion continue to show a status of disabled.

Note If all health modules or a device are added to the exclusion list, the device's health status will be set to disabled. This status refers only to health monitoring and should not be interpreted as the device being offline or disconnected.

You can also disable an individual health module. For example, when you reach the host limit on the Cloud-Delivered Firewall Management Center, you can disable Host Limit status messages. Excluding health modules for individual interfaces is not supported on devices operating in transparent mode.

Note that on the main Health Monitor page you can distinguish between devices that are excluded if you expand to view the list of devices with a particular status by clicking the arrow in that status row.

Note On Cloud-Delivered Firewall Management Center, Health Monitor exclusion settings are local configuration settings. Therefore, if you exclude a device, then delete it and later re-register it with the Cloud-Delivered Firewall Management Center, the exclusion settings remain persistent. The newly re-registered device remains excluded.

Temporary device exclusion scenario

You can use the health monitor exclude feature to disable health monitoring status reporting for a device, Cloud-Delivered Firewall Management Center where applicable, or specific health modules. For example, if you know that a segment of your network will be unavailable, you can temporarily disable health monitoring for a managed device on that segment to prevent the health status on the Cloud-Delivered Firewall Management Center from displaying a warning or critical state because of the lapsed connection to the device.

To temporarily disable health events from a managed device, or from Cloud-Delivered Firewall Management Center where supported, go to the exclusion configuration page and add the target to the exclusion list. After the setting takes effect, the system no longer considers the excluded device when calculating the overall health status. The Health Monitor Appliance Status Summary lists the device as disabled.

Exclude devices from health monitoring

Exclude devices from health monitoring to disable their events and health status when needed for maintenance or troubleshooting purposes.

You can exclude managed devices individually or by group, model, associated health policy, or other supported scope. In Cloud-Delivered Firewall Management Center exclusion can also apply to Cloud-Delivered Firewall Management Center where supported.

If you need to set the events and health status for an individual managed device to disabled, you can exclude the device from health monitoring. Where supported, you can also exclude Cloud-Delivered Firewall Management Center health monitoring. After the exclusion settings take effect, the device shows as disabled in the Health Monitor Appliance Module Summary, and health events for the device have a status of disabled.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Exclude**.

Step 2 Click **Add Device**.

Step 3 In the **Device Exclusion** dialog box, under **Available Devices**, click **Add** (⊕) against the device that you want to exclude from health monitoring.

Step 4 Click **Exclude**.

The selected device is displayed in the exclusion main page.

Step 5 To remove the device from the exclusion list, click **Delete** (🗑).

Step 6 Click **Apply**.

Exclude health policy modules on a device

Excluding health policy modules is useful when certain modules produce irrelevant or temporary health alerts, which you do not want to affect the overall device status. Exclusion settings indicate the number of modules excluded from health monitoring.

Note Excluding health modules for individual interfaces is not supported on devices operating in transparent mode.

Tip Make sure that you keep track of individually excluded modules so you can reactivate them when you need them. You may miss essential warning or critical messages if you accidentally leave a module disabled.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Exclude**.

Step 2 Click **Edit** (✎) next to the Firewall Threat Defense device you want to modify.

Step 3 In the **Exclude Health Modules** dialog box, by default, all the modules of the device are excluded from health monitoring.

Certain modules are applicable to a specific device only; for more information, refer to Health modules, on page 373.

Step 4 To choose modules to be excluded from health monitoring, click the **Enable Module Level Exclusion** link.

The **Exclude Health Modules** dialog box displays all the modules of the device. The modules that are not applicable for the associated health policies are disabled by default. To exclude a module, perform these steps:

a. Click the **Slider** (⬭) button next to the desired module.

b. To specify the duration of the exclusion for the selected modules, from the **Exclude Period** drop-down list, select the duration.

Step 5 (Optional) To stop receiving health status updates for physical interfaces and instead focus on the health status of their subinterfaces, you can disable health monitoring for the physical interfaces while continuing to monitor and receive health alerts for the subinterfaces.

a) If not already enabled, click the toggle button next to the **Interface statistics** module.
b) From the **Exclude Period** drop-down list, choose a duration for which the health status of the physical interfaces is not required.
c) Click the **Exclude specific interfaces** radio button.
d) Check the check boxes next to the physical interfaces that you want to exclude.
e) Selecting a physical interface will automatically exclude both the physical interface and its subinterfaces. Therefore, uncheck the **Subinterfaces** check box next to the corresponding physical interface to continue receiving health alerts for the subinterfaces.

Step 6 If you select an **Exclude Period** other than **Permanent**, for your exclusion configuration, you can choose to automatically delete the configuration when it expires. To enable this setting, check the **Auto-delete expired configurations** check box.

Step 7 Click **OK**.

Step 8 Click **Apply**.

Note
Interface exclusion differs from other module exclusions because you can exclude selected physical interfaces while still monitoring associated subinterfaces, depending on the UI options available in your version.

Exclude interface health monitoring

Use interface exclusion when you want to suppress health status updates for selected physical interfaces without suppressing health monitoring for the entire device. Interface exclusion is useful when interface state is expected during maintenance or when you want to focus on subinterface health instead of physical interface status.

When you configure interface exclusion, select the physical interfaces that you want to exclude from the list of available interfaces. Depending on the available UI options in your version, selecting a physical interface can also exclude its subinterfaces. If subinterface controls are available, you can continue monitoring subinterfaces while excluding only the physical interface.

Use this option to reduce unnecessary health alerts when traffic loss or interface state changes are expected on a physical interface.

Manage expired health monitor exclusions

When monitoring device or module health, some items may have temporary exclusions. Once these exclusion periods lapse, you need to decide whether to renew the exclusion or clear the device/module from exclusion so alerts resume.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Exclude**.

The **Warning** (⚠) icon is displayed against the device indicating the expiry of the duration of exclusion of the device or the modules from alerting.

Step 2 To renew the exclusion of the device, click **Edit** (✎) next to the device. In the **Exclude Health Modules** dialog box, click the **Renew** link. The exclusion period of the device is extended with the current value.

Step 3 To clear the device from being excluded, click **Delete** (🗑) next to the device, click **Remove the device from exclusion**, and then click **Apply**.

Step 4 To renew or clear the modules from exclusion, click **Edit** (✎) next to the device. In the **Exclude Health Modules** dialog box, click the **Enable Module Level Exclusion** link, and then click the **Renew** or **Clear** link

against the modules. When you click **Renew**, the exclusion period is extended on the module with the current value.

Health monitor alerts

A health monitor alert is a notification mechanism that

- provides email or webhook notifications when module status changes occur,
- enables association of alert responses with health event levels to trigger notifications based on event severity, and
- allows customization to monitor specific system conditions such as disk space thresholds.

Disk space monitoring alert example

If you are concerned that your appliances may run out of hard disk space, you can automatically send an email to a system administrator when the remaining disk space reaches the warning level. If the hard drive continues to fill, you can send a second email when the hard drive reaches the critical level.

Health monitor alert information

The alerts generated by the health monitor contain specific information that helps identify issues and their severity.

Health monitor alerts include these attributes:

- Severity: Indicates the severity level of the alert.
- Module: Specifies the health module whose test results triggered the alert.
- Description: Includes the health test results that triggered the alert.

The table describes these severity levels.

Table 21: Alert severities

Severity	Description
Critical	The health test results met the criteria to trigger a Critical alert status.
Warning	The health test results met the criteria to trigger a Warning alert status.
Normal	The health test results met the criteria to trigger a Normal alert status.
Error	The health test did not run.
Recovered	The health test results met the criteria to return to a normal alert status, following a Critical or Warning alert status.

Health alerts for Secure Firewall 200 Series device

To optimize performance and ensure effective resource utilization, the health alerts in the Secure Firewall 200 Series device are limited only to the essential health modules. This table lists the health modules in the Secure Firewall 200 Series device, which generate health alerts. You can view all the metrics in the health monitoring dashboard, which is similar to that of other Threat Defense device models.

Table 22: Health Alerts for Secure Firewall 200 Series Device

Health Module	Health Alert
Certificate Monitoring	Alerts when service authentication certificates are nearing expiration or have expired.
Cluster/HA Failure Status	Provide alerts when a device joins, leaves, or is elected as the primary unit.
Database	Provide alerts on database integrity issues related to schema or configuration data.
Disk Usage	Monitors disk usage in the device's hard drive and alerts when usage exceeds the configured thresholds.
Disk Status	Provide alerts on hard disk or RAID controller issues.
Firewall Threat Defense Platform Faults	Monitors platform faults and generates health alerts for them.
FXOS Health	Alerts when the FXOS HTTPS service is not running in the device.
Identity Process	Monitors the health and operation of identity-related services.
Inline Link Mismatch Alarms	Provide alerts if inline pair interfaces negotiate different speeds.
Interface Statistics	Determines if the device currently collects traffic and alerts based on the traffic status of physical interfaces and aggregate interfaces.
Out of band Configuration Changes	Alerts when there is a conflict between the existing Cloud-Delivered Firewall Management Center configuration and the out-of-band configuration changes that are made.
Process Status	Provide alerts when processes in the device are terminated outside of the process manager.
Snort Identity Memory Usage	Enables you to set a warning threshold for Snort identity processing, and alerts when memory usage exceeds the level that is configured for the module.
Snort Reconfiguring Detection	Alerts if a device reconfiguration has failed.

Health Module	Health Alert
Threat Data Updates on Devices	Monitors updates of threat intelligence data and alerts if this information has not been updated in the devices within the time period you have specified.

Caution Generating troubleshooting files in a Firewall Threat Defense device is a CPU-intensive task. Because of limited CPU resources in the Secure Firewall 200 Series device, you may observe higher CPU usage and associated health alerts during this process. To prevent any potential traffic disruption, it's recommended to generate troubleshooting files only when the device is not actively handling network traffic.

Note

- The Secure Firewall 200 Series device raises alerts only for database integrity issues-related configuration data (sometimes called EO).
- The Secure Firewall 200 Series device does not maintain a local URL database, and supports cloud-only URL lookups. Local URL database-related alerts are not available for this device type.

Health event external notifications

Firewall Management Center uses separate configurations to detect a health condition and notify an external system about it.

A health policy defines what Cloud-Delivered Firewall Management Center monitors and which conditions produce a health event. An alert response configuration defines how Cloud-Delivered Firewall Management Center communicates with an external destination, such as an email server. The monitor alert rule configuration connects the health policy and alert response by specifying which health events should trigger a particular alert response.

How the configurations work together

These steps describe how you configure external notifications for health events and how Cloud-Delivered Firewall Management Center processes the health event.

1. Configure and apply the health policy: A health policy enables health modules and defines the criteria used to assess appliance health.

   For example, a health policy can enable the **Memory Usage** module and define **Critical** thresholds.

   Apply the health policy to the appliance, including the Cloud-Delivered Firewall Management Center when you want to monitor Cloud-Delivered Firewall Management Center resources.

2. Configure the alert response: An alert response defines how and where Cloud-Delivered Firewall Management Center sends the external notification. Depending on the response type, the alert response contains the destination address, server or endpoint information, and other communication settings.

   For example, an email alert response identifies the recipients, sender, and mail relay host that Cloud-Delivered Firewall Management Center uses to send the message.

For more information about alert responses, refer to Configuring external alerts with alert responses, on page 487.

3. Configure the monitor alert rule: A monitor alert rule associates a health-event severity and one or more health modules with an alert response.

   For example, a health monitor alert rule can specify:

   - Severity: **Critical**
   - Health module: **Memory Usage**
   - Alert response: **Critical memory email**

   This means that when Cloud-Delivered Firewall Management Center generates a **Critical** event for the **Memory Usage** module, it invokes the associated alert response.

   For more information about health monitor alerts, refer toHealth monitor alerts, on page 393.

4. Cloud-Delivered Firewall Management Center generates a health event: Cloud-Delivered Firewall Management Center evaluates the enabled health modules according to the health policy and its configured run interval. When a health module evaluates an appliance and the result meets a configured condition, Cloud-Delivered Firewall Management Center generates a health event.

   Note that a health event is an event log within Cloud-Delivered Firewall Management Center; it does not send an external notification by itself.

5. Cloud-Delivered Firewall Management Center matches the event and sends the notification: Cloud-Delivered Firewall Management Center evaluates the generated event against the configured health monitor alerts. If the event's module and severity match a monitor alert rule, Cloud-Delivered Firewall Management Center invokes the alert response associated with that rule and sends the notification to the configured external system.



Note Health alerts can take approximately 5–6 minutes to be generated after the corresponding health event occurs.

Create a health monitor alert

When you create a health monitor alert, you create an association between a severity level, a health module, and an alert response. You can use an existing alert or configure a new one specifically to report on system health. When the severity level occurs for the selected module, the alert triggers.

If you create or update a threshold in a way that duplicates an existing threshold, you are notified of the conflict. When duplicate thresholds exist, the health monitor uses the threshold that generates the fewest alerts and ignores the others. The timeout value for the threshold must be between 5 and 4,294,967,295 minutes.

Before you begin

- Ensure that you have configured alert responses that govern the Cloud-Delivered Firewall Management Center's communication with the email server or webhook where you send the health alert; For more information, refer to Configuring external alerts with alert responses, on page 487.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor Alerts**.

Step 2 Click **Add**.

Step 3 In the **Add Health Alert** dialog box, enter a name for the health alert in the **Health Alert Name** field.

Step 4 From the **Severity** drop-down list, choose the severity level you want to use to trigger the alert.

Step 5 From the **Alert** drop-down list, choose the alert response that you want to trigger when the specified severity level is reached. If you have not yet configured the alert responses, click **Alerts** to visit the **Alerts** page and set them.

Step 6 From the **Health Modules** list, choose the health policy modules for which you want the alert to apply.

Step 7 Optionally, in the **Threshold Timeout** field, enter the number of minutes that should elapse before each threshold period ends and the threshold count resets.

Even if the policy run time interval value is less than the threshold timeout value, the interval between two reported health events from a given module is always greater. For example, if you change the threshold timeout to 8 minutes and the policy run time interval is 5 minutes, there is a 10-minute interval (5 x 2) between reported events.

Step 8 Click **Save** to save the health alert.

Edit a health monitor alert

Administrators may need to update the alert parameters for health monitor alerts in order to adjust severity, change modules, or select different responses as requirements evolve.

Before you begin

Follow these steps to edit health monitor alerts:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor Alerts**.

Step 2 Click the **Edit** (✎) icon that is provided against the required health alert that you want to modify.

Step 3 In the **Edit Health Alert** dialog box, from the **Alert** drop-down list, select the required alert entry, or click **Alerts** link to configure a new alert entry.

Step 4 Click **Save**.

Delete health monitor alerts

Health monitor alerts notify you about system events, issues, or warnings. You may need to delete outdated or irrelevant alerts to streamline monitoring.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor Alerts**.

Step 2 Click **Delete** (🗑) next to the health alert you want to delete.

Step 3 Click **Delete health alert** to delete it.

What to do next

Disable or delete the underlying alert response to ensure that alerting does not continue. For more information, refer to Configuring external alerts with alert responses, on page 487.

About the health monitor

The health monitor provides the compiled health status for all devices managed by the Cloud-Delivered Firewall Management Center, plus the Cloud-Delivered Firewall Management Center itself.

The health monitor is composed of:

- The Health Status summary page — Provides you with an at-a-glance view of the health of the Cloud-Delivered Firewall Management Center and all of the devices that the Cloud-Delivered Firewall Management Center manages. In a multi-domain deployment, you can view the health status summary for a device in the domain where the device is located. Devices are listed individually, or grouped according to their geolocation, high availability, or cluster status where applicable.
  - View the health summary of the Cloud-Delivered Firewall Management Center and any device when you hover on the hexagon that represents the device health.
  - The dot to the left of a device indicates its health:
    - Green — No alarms.
    - Orange — At least one health warning.
    - Red — At least one critical health alarm.
- The Monitoring navigation pane — Allows you to navigate the device hierarchy. You can view health monitors for individual devices from the navigation pane.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 View the status of the Cloud-Delivered Firewall Management Center and its managed devices in the **Health Status** landing page.

a) Hover your pointer over a hexagon to view the health summary of a device. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

b) In the device list, click **Expand**( ) and **Collapse** ( ) to expand and collapse the list of health alerts for a device.

When you expand the row, all of the health alerts are listed, including the status, title, and details.

Note
Health alerts are sorted by their severity level.

Step 3 Use the Monitoring navigation pane to access device-specific health monitors.

a) In the device list, click **Expand** ( ) and **Collapse** ( ) to expand and collapse the list of managed devices.

When you expand the row, all of the devices are listed.

b) Click on a device to view a device-specific health monitor.

c) In the health monitor, hover over a graph to view all metrics and their respective values at that specific point on the graph. Click on the graph to pin the metrics statistics box, allowing you to explore the metrics and their values in detail. To close the statistics box and move to another point on the graph, simply click the close button.

You have accessed the health monitor and can view the compiled health status for all managed devices. You can interact with the Health Status landing page and use the Monitoring navigation pane to access device-specific health information.

What to do next

- See Device health monitors, on page 399 for information about the compiled health status and metrics for any device managed by the Cloud-Delivered Firewall Management Center.

Device health monitors

The device health monitor provides the compiled health status for any device managed by the Cloud-Delivered Firewall Management Center. The device health monitor collects health metrics for Secure Firewall devices in order to predict and respond to system events.

The device health monitor is comprised of these components:

- System Details: Displays information about the managed device, including the installed Secure Firewall version and other deployment details.
- Troubleshooting & Links: Provides convenient links to frequently used troubleshooting topics and procedures.
- Health alerts: A health alert monitor provides an at-a-glance view of the health of the device.
- Time range: An adjustable time window to constrain the information that appears in the various device metrics windows.

- Device metrics: An array of key firewall device health metrics categorized across predefined dashboards, including:
  - CPU: CPU utilization, including the CPU usage by process and by physical cores. The Firewall Threat Defense CPU core allocation dashboard shows core assignments for these categories:
    - **Data Plane**: Handles basic network functions, including core packet forwarding and network data processing.
    - **Snort**: Manages intrusion detection and deep packet inspection features.

Note Multi-threaded Snort processes are allocated across multiple Snort cores to improve device performance. The Snort core allocation shown in the **CPU** dashboard reflects the total number of Snort process threads assigned to Snort cores.

    - **System**: Includes all other system processes. While some processes may have dedicated CPU cores allocated, their usage is combined and displayed under the **System** category in the dashboard.
  - Memory: Device memory utilization, including data plane and Snort memory usage.
  - Interfaces: Interface status and aggregate traffic statistics.
  - Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
  - Snort: Statistics related to the Snort process.
  - Disk Usage: Device disk usage, including the disk size and disk utilization per partition.
  - Critical Processes: Statistics related to managed processes, including process restarts and other select health monitors such as CPU and memory utilization.

Note During a device upgrade or high-availability failover event, the Firewall Threat Defense device may briefly appear as **Offline** in the device's health monitoring dashboard. This happens because health alerts are cleared during the process and are only updated after the process is complete. Wait for the upgrade or failover operation to finish.

For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

View system details and troubleshoot devices

Use this task when you need to monitor device health or initiate troubleshooting procedures on specific devices through the management center.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed devices.

Step 3 Click on a device to view a device-specific health monitor.

Step 4 Click the link for **View System & Troubleshoot Details …**

This panel is collapsed by default. Clicking on the link expands the collapsed section to see **System Details** and **Troubleshooting & Links** for the device. The system details include:

- **Version:** The Secure Firewall software version.
- **Model:** The device model.
- **Mode:** The firewall mode. The Firewall Threat Defense device supports two firewall modes for regular firewall interfaces: Routed mode and Transparent mode.
- **VDB:** The Cisco vulnerability database (VDB) version.
- **SRU:** The intrusion rule set version.
- **Snort:** The Snort version.

Step 5 You have these troubleshoot choices:

- Generate troubleshooting files; refer to Generate troubleshooting files for specific system functions, on page 433
- Generate and download advanced troubleshooting files; refer to Download advanced troubleshooting files, on page 434.
- Create and modify health policies; refer to Create a health policy, on page 385.
- Create and modify health monitor alerts; refer to Create a health monitor alert, on page 396.

View the device health monitor

The device health monitor compiles device metrics and displays health status and trends using an array of dashboards.

Follow these steps to view the device health monitor.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 2 In the device list, click **Expand**( ) and **Collapse** ( ) to expand and collapse the list of managed devices.

Step 3 View the **Health Alerts** for the device in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the device. The pop-up window shows a truncated summary of the top five health alerts. Click the pop-up to open a detailed view of the health alert summary.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click the **Show the deployment details on top of the graph** ( ) icon for a deployment overlay on the trend graph, with respect to the selected time range.

The icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. In the case of multiple deployments, multiple bands/lines can appear. Click the icon on top of the dotted line to view the deployment details.

Step 6 View the device monitor reports health and performance metrics in several predefined dashboards by default.

The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.
- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics related to the Snort process.
- ASP Drops: Statistics related to the Accelerated Security Path (ASP) performance and behavior.

Note

- Multi-threaded Snort processes are allocated across multiple Snort cores to improve device performance. Note that the Snort core allocation shown in the **CPU**dashboard reflects the total number of Snort process threads assigned to Snort cores.
- The **Process Health** widget in Cloud-Delivered Firewall Management Center displays the CPU% usage for each process. The CPU% metric reflects usage relative to the number of cores in the Firewall Threat Defense device, where 100% corresponds to full utilization of one core. For example, on an 8-core device, 200% indicates two fully utilized cores, leaving six available for other processes. To monitor overall system CPU usage, use the **CPU** widget instead of the **Process Health** widget.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 7 Click the **Add New Dashboard**(+) to create a custom correlation dashboard by building your own variable set from the available metric groups.

For more information, refer to Correlating device metrics, on page 403.

Correlating device metrics

The device health monitor includes an array of key Firewall Threat Defense device metrics that serve to predict and respond to system events. The health of any Firewall Threat Defense device can be determined by these reported metrics.

The device monitor reports these metrics in several predefined dashboards by default. These dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.
- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics related to the Snort process.
- ASP Drops: Statistics related to the Accelerated Security Path (ASP) performance and behavior.

You can add custom dashboards to correlate metrics that are interrelated. Select from predefined correlation groups, such as CPU and Snort; or create a custom correlation dashboard by building your own variable set from the available metric groups. For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Before you begin

- To view and correlate the time series data (device metrics) in the health monitor dashboard, enable REST API (**System** (⚙) > **Configuration** > **REST API Preferences**).

Note Correlating device metrics is available only for Firewall Threat Defense 6.7 and later versions. Hence, for Firewall Threat Defense versions earlier than 6.7, the health monitor dashboard does not display these metrics even if you enable REST API.

Follow these steps to correlate device metrics:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**. In the **Devices** list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed devices, and choose the device for which you want to modify the dashboard.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 2 Click the **Add New Dashboard**(+) icon to add a new dashboard, and specify a name to identify the dashboard.

Step 3 To create a dashboard from a predefined correlation group, click **Add from Predefined Correlations** drop-down, choose the group, and click **Add Dashboard**. To create a custom correlation dashboard, choose a group from the **Select Metric Group** drop-down, then choose corresponding metrics from the **Select Metrics** drop-down.

For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 Click **Add Metrics** to add and select metrics from another group, and to remove an individual metric, click the Remove ✕ icon on the right side of the item. Click the delete icon to remove the entire group. Click **Add Dashboard** to add the dashboard to the health monitor.

Step 5 You can **Edit** or **Delete** the predefined dashboards and the custom correlation dashboards.

Cluster health monitor

When a Firewall Threat Defense is the control node of a cluster, the Cloud-Delivered Firewall Management Center collects various metrics periodically from the device metric data collector.

The cluster health monitor is comprised of these components:

- Overview dashboard: Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Cloud-Delivered Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.

Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard: Displays load distribution across the cluster nodes in two widgets:

- The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
- The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.

- Member Performance dashboard: Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard: Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links: Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range: An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard: Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

View the cluster health monitor

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

Ensure you have created a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Follow these steps to view the cluster health monitor:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed cluster devices.

Step 3 Click on the cluster name to view the cluster health statistics.

The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution: Traffic and packet distribution across the cluster nodes.

- Member Performance: Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL: Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 View device monitor health and performance metrics in several predefined dashboards.

The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.
- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics that are related to the Snort process.
- ASP drops: Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. Refer to Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Health monitor status categories

The table summarizes the available health monitor status categories, their severity levels, icons, associated colors in pie charts, and descriptions.

Table 23: Health status indicator

Status Level	Status Icon	Status Color in Pie Chart	Description
Error	**Error** ()	Black	Indicates that at least one health monitoring module has failed on the appliance and has not been successfully re-run since the failure occurred. Contact your technical support representative to obtain an update to the health monitoring module.
Critical	**Critical** ()	Red	Indicates that the critical limits have been exceeded for at least one health module on the appliance and the problem has not been corrected.
Warning	**Warning** ()	Yellow	Indicates that warning limits have been exceeded for at least one health module on the appliance and the problem has not been corrected. This status also indicates a transitionary state, where, the required data is temporarily unavailable or could not be processed because of changes in the device configuration. Depending on the monitoring cycle, this transitionary state is auto-corrected.
Normal	**Normal** ()	Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance.
Recovered	**Recovered** ()	Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance, including modules that were in a Critical or Warning state.
Disabled	**Disabled** ()	Blue	Indicates that an appliance is disabled or excluded, that the appliance does not have a health policy applied to it, or that the appliance is currently unreachable.

Health event views

Health event views enable you to monitor health events logged by the health monitor in the Cloud-Delivered Firewall Management Center. The views are fully customizable, allowing you to analyze health statuses efficiently and access related event data.

You can search event data to easily access other information that may be related to the events you are investigating. If you understand what conditions each health module tests for, you can more effectively configure alerting for health events.

You can perform many of the standard event view functions on the health event view pages.

View health events

Use this task when you need to monitor the health of your appliances and review any incidents or events recorded.

When you access health events from the Health Monitor page on your Cloud-Delivered Firewall Management Center, you retrieve all health events for all managed appliances.

Tip You can bookmark this view to allow you to return to the page in the health events workflow containing the Health Events table of events. The bookmarked view retrieves events within the time range you are currently viewing, but you can then modify the time range to update the table with more recent information if needed.

Follow these steps to view health events.

Procedure

Choose **System** (⚙) > **Health** > **Events**.

Tip
If you are using a custom workflow that does not include the table view of health events, click **(switch workflow)**. On the **Select Workflow** page, click **Health Events**.

Note
If no events appear, you may need to adjust the time range.

View the health events table

The Health Events Table provides actionable insight into system health and allows users to organize, modify, and respond to health events.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Events**.

Step 2 You have the following choices:

- Bookmark: To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**, provide a name for the bookmark, and click **Save**.

- Change Workflow: To choose another health events workflow, click **(switch workflow)**.
- Delete Events: To delete health events, check the check box next to the events you want to delete, and click **Delete**. To delete all the events in the current constrained view, click **Delete All**, then confirm you want to delete all the events.
- Generate Reports: Generate a report based on data in the table view — click **Report Designer**.
- Modify: Modify the time and date range for events listed in the Health table view. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
- Navigate: Navigate through event view pages.
- Navigate Bookmark: To navigate to the bookmark management page, click **View Bookmarks** from any event view.
- Navigate Other: Navigate to other event tables to view associated events.
- Sort: Sort the events that appear, change what columns display in the table of events, or constrain the events that appear
- View All: To view event details for all events in the view, click **View All**.
- View Details: To view the details associated with a single health event, click the down arrow link on the left side of the event.
- View Multiple: To view event details for multiple health events, choose the check box next to the rows that correspond with the events you want to view details for and then click **View**.
- View Status: To view all events of a particular status, click status in the Status column for an event with that status.
- Search: Search for specific health events or health metrics recorded within your deployment. Customize your search query to analyze health events and other monitored data and save the search constraints for quick access. For more information, refer to Event searches, on page 519.

Health event fields

The health monitor modules you enable in your health policy run tests to determine appliance health status. When the health status meets specified criteria, a health event is generated. The health event table includes fields that are searchable and viewable for monitoring appliance health.

Table 24: Health event fields

Field	Description
Module Name	Specify the name of the module which generated the health events you want to view. For example, to view events that measure CPU performance, type **CPU**. The search should retrieve applicable CPU Usage and CPU temperature events.
Test Name (Search only)	The name of the health module that generated the event.
Time (Search only)	The timestamp for the health event.

Field	Description
Description	The description of the health module that generated the event. For example, health events generated when a process was unable to execute are labeled **Unable to Execute**.
Value	The value (number of units) of the result obtained by the health test that generated the event. For example, if the Cloud-Delivered Firewall Management Center generates a health event whenever a device it is monitoring is using 80 percent or more of its CPU resources, the value could be a number from 80 to 100.
Units	The units descriptor for the result. You can use the asterisk (*) to create wildcard searches. For example, if the Cloud-Delivered Firewall Management Center generates a health event when a device it is monitoring is using 80 percent or more of its CPU resources, the units descriptor is a percentage sign (%).
Status	The status (Error, Critical, Warning, Normal, Recovered, or Disabled) reported for the appliance.
Device	The appliance where the health event was reported.

Sending device health metrics to OpenConfig

OpenConfig is a vendor-independent software layer that provides a single way of streaming network telemetry data to multiple vendors to manage and monitor networks. The OpenConfig streaming telemetry option in the secure firewall uses gNMI (gRPC Network Management Interface) protocol and allows you to control and generate telemetry streams from your Firewall Threat Defense devices to a data collection system.

The Firewall Threat Defense health policy contains all the configurations to support and enable the OpenConfig streaming telemetry functionality. When you deploy the health policy to the device, the OpenConfig streaming telemetry configuration activates a gNMI server and starts listening to Remote Procedure Call (RPC) messages from the data collectors.

Subscription model of OpenConfig telemetry streams

OpenConfig uses a subscription-based model where the data collectors query the Firewall Threat Defense devices for telemetry data or act as collector for the streamed telemetry data. When a data collector wishes to receive updates and metrics from the Firewall Threat Defense device, it sends a subscribeRequest RPC message to the Firewall Threat Defense gNMI server. The subscription request includes details of one or more paths to which the data collector wishes to subscribe. The message also includes subscription mode which describes the longevity of the subscription.

The Firewall Threat Defense server supports these subscription modes:

- *Once subscription*: The Firewall Threat Defense device sends requested data to the gNMI paths only once.
- *Stream subscription*: The Firewall Threat Defense continuously streams telemetry data according to the triggers specified in the `SubscribeRequest` RPC message.

- *Sampled subscription*: The Firewall Threat Defense server streams the requested data as per the interval specified in the subscription message. The minimum interval that the threat defense support is one minute.
- *On-change subscription*: The Firewall Threat Defense sends the data whenever the requested values change.

The Firewall Threat Defense server generates `SubscribeResponse` RPC messages according to the type of subscription that is created, at the frequency requested by the data collectors.

You can use these deployment modes for OpenConfig streaming telemetry configuration:

- **DIAL-IN**: In this mode, the gNMI server opens a port on the Firewall Threat Defense and waits for `SubscribeRequest` RPC messages from data collectors. In the device health policy, you can specify the port number to use by the gNMI server and the IP address of the data collector that can connect with the gNMI service. If not specified, the gNMI server uses port number 50051. The DIAL-in mode is ideal to use in a trusted network where the endpoint that subscribes to telemetry streams are trusted.
- **DIAL-OUT**: The gNMI service is designed to work in server mode where it accepts subscription requests from gNMI data collectors and serve the telemetry data. If the gNMI data collectors cannot reach the gNMI server, the Firewall Threat Defense uses a tunnel client and establishes a gRPC tunnel with the external server. This tunnel allows exchange of RPC messages between gNMI server and client. The DIAL-OUT mode is ideal to use when the data collectors are hosted on the cloud or outside the trusted network.

In both DIAL-IN and DIAL-OUT mode, all the communication between gNMI server and gNMI client uses TLS encryption and this requires to generate a set of certificates with private keys for the TLS encryption. DIAL-OUT mode requires extra keys for the tunnel infrastructure. For more information, refer to How to Generate Certificate with Private Key.

Generate certificates and private keys

Certificates and private keys are required for OpenConfig streaming telemetry. You must use the same certificate authority (CA) for all components to ensure secure and consistent authentication.

Note To ensure that you generate certificates using the same CA, run these commands together and from the same endpoint. If you want to retry the commands, you must retry all commands.

Procedure

Step 1 Make a folder, for example `keys`, in the endpoint where you want to run these commands.

Example:

```
mkdir keys
```

Step 2 Create a self-signed CA certificate with a corresponding private key.

Example:

This sample command generates a new RSA private key and uses it to create a self-signed X.509 certificate with provided subject information:

```
openssl req -x509 -newkey rsa:4096 -days 365 -nodes -keyout keys/ca-key.pem -out
keys/ca-cert.pem -subj "/C=XX
/ST=YY/L=ZZZ/O=Example/OU=EN/CN=gnmi-ca/emailAddress=abc@example.com"
```

The subject information includes the provided Country (C), State (ST), Locality (L), Organization (O), Organizational Unit (OU), Common Name (CN), and email address.

The private key is saved as `CA-key.pem` file, and the certificate is saved as `CA-cert.pem` file in the `keys` folder.

Step 3 Create a self-signed server certificate with the specified Common Name (CN) and Subject Alternative Name (SAN):

Example:

This sample command generates a new RSA private key and uses it to create a self-signed X.509 certificate with provided subject information. In this example, 192.168.0.200 is the IP address of the Firewall Threat Defense device and 192.168.0.202 is the IP address of the client.

Note
Client IP is not required if you want to use this certificate and key sets in dial-in mode.

```
CN="192.168.0.200"
SAN="IP:192.168.0.200,IP:192.168.0.202"
openssl req -newkey rsa:4096 -nodes -keyout keys/server-key.pem -out keys/server-req.pem
-subj "/C=XX/ST=YY/L=ZZZ/O=Example/OU=EN/CN=${CN}/emailAddress=abc@example.com}"
openssl x509 -req -extfile <(printf "subjectAltName=${SAN}") -in keys/server-req.pem -days
 60 -CA keys/ca-cert.pem -CAkey keys/ca-key.pem -CAcreateserial -out keys/server-cert.pem
cat keys/server-key.pem keys/server-cert.pem keys/ca-cert.pem > keys/server-combined.pem
```

The `openssl req` command generates a new RSA private key and a Certificate Signing Request (CSR). The private key is saved as `server-key.pem` file, and the CSR is saved as `server-req.pem` file in the `keys` folder.

The `openssl x509` command processes the CSR and generates a server certificate. The server certificate is saved as `server-cert.pem` file in the `keys` folder.

The `cat` command combines the server key, server certificate, and the CA certificate into a single file named `server-combined.pem` and saves the file in the `keys` folder.

You have to upload the `server-combined.pem` while configuring **OpenConfig Streaming telemetry** from the Cloud-Delivered Firewall Management Center. The gNMI server that runs on the Firewall Threat Defense and the tunnel server (dial-out mode) uses this certificate for TLS communication. If you encrypt the private key with a passphrase, ensure that you specify the passphrase while uploading the certificate to the Cloud-Delivered Firewall Management Center.

Step 4 Create client certificate with the specified Common Name (CN) and Subject Alternative Name (SAN).

Example:

This sample command generates a new RSA private key and uses it to create a self-signed X.509 certificate with provided subject information. In this example, 192.168.0.202 is the IP address of the client.

```
CN="192.168.0.202"
SAN="IP:192.168.0.202"
openssl req -newkey rsa:4096 -nodes -keyout keys/client-key.pem -out keys/client-req.pem
-subj "/C=XX/ST=YY/L=ZZZ/O=example/OU=EN/CN=${CN}/emailAddress=abc@example.com"
openssl x509 -req -extfile <(printf "subjectAltName=${SAN}") -in keys/client-req.pem -days
 60 -CA keys/ca-cert.pem -CAkey keys/ca-key.pem -CAcreateserial -out keys/client-cert.pem
```

The gNMI client uses the client certificate `client-cert.pem` and the private key for TLS communication.

Step 5 (Optional) For dial-out mode, create the tunnel server certificate with the specified Common Name (CN) and Subject Alternative Name (SAN).

Example:

This sample command generates a new RSA private key and uses it to create a self-signed X.509 certificate with provided subject information. In this example, 192.168.0.202 is the IP address of the client.

```
CN="192.168.0.202"
SAN="IP:192.168.0.202"
openssl req -newkey rsa:4096 -nodes -keyout keys/tunnel-server-key.pem -out
keys/tunnel-server-req.pem -subj "
/C=XX/ST=YY/L=ZZZ/O=Example/OU=EN/CN=${CN}/emailAddress=abc@example.com}"
openssl x509 -req -extfile <(printf "subjectAltName=${SAN}") -in keys/tunnel-server-req.pem
 -days 60 -CA keys/ca-cert.pem -CAkey keys/ca-key.pem -CAcreateserial -out
keys/tunnel-server-cert.pem
```

Configure OpenConfig streaming telemetry

OpenConfig streaming telemetry allows network monitoring systems to collect real-time data from threat defense devices using the gNMI protocol. This configuration supports both DIAL-IN and DIAL-OUT deployment modes for flexible integration with existing network management infrastructure.

Before you begin

- Make sure the Firewall Threat Defense device where you want to deploy the health policy configuration allows installation of the SSL certificate and PRIVATE KEY.
- Ensure that you configure a gNMI client that supports the OpenConfig streaming telemetry implementation, from which you can make the gRPC requests to the gNMI server on the Firewall Threat Defense.
- To use DIAL-OUT mode and configure OpenConfig streaming telemetry, ensure that you configure a gRPC tunnel server and client on the management system. This tunnel configuration enables communication between the gNMI client and the Firewall Threat Defense device.

Follow these steps to configure OpenConfig streaming telemetry:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Policy**.

Step 2 Click the **Edit health policy** icon next to the threat defense health policy that you want to modify.

Step 3 Go to the **Settings** tab.

Step 4 Move the **OpenConfig Streaming Telemetry** slider to enable the configuration. This configuration is disabled by default.

Step 5 Upload the **SSL Certificate**. The gNMI server uses this certificate to enable server authentication for the TLS connection and encrypt all communications through the channel.

The OpenConfig streaming telemetry configuration supports only certificate with PEM format.

Step 6 (Optional) Specify the Passphrase if the PRIVATE KEY files are ENCRYPTED.

Step 7 Choose the deployment mode to use for streaming telemetry over gNMI protocol.

For **DIAL-IN** mode:

a. Assign a port number for the gNMI service.

The gNMI server opens the port and waits for gRPC requests from the collector.

b. Specify the IPv4/IPv6 address of the gNMI collectors that can connect to the Firewall Threat Defense device.

c. Click **Add Collector** to add more gNMI collectors. You can add a maximum of five collectors.

For **DIAL-OUT** mode:

a. Specify the hostname and port number for the gNMI collector, which can subscribe to streaming telemetry from the Firewall Threat Defense device.

b. Click **Add Collector** to add more gNMI collectors. You can add a maximum of five collectors.

Step 8 Specify the username and password to validate the gNMI collector.

The Firewall Threat Defense server uses this credential to authenticate the gNMI collector when receiving the **SubscribeRequest** RPC message. Each telemetry message is not authenticated using the username and password. The system uses the previously authenticated ENCRYPTED streaming channel to carry telemetry messages.

Step 9 Click **Save**.

What to do next

Deploy the health policy to your Firewall Threat Defense device, for the configuration changes to take effect.

Troubleshoot OpenConfig streaming telemetry

This reference provides troubleshooting procedures for OpenConfig streaming telemetry issues including certificate errors, server response failures, and status verification methods.

Certificate signed by unknown authority

- Ensure that you have uploaded the correct certificate to the Cloud-Delivered Firewall Management Center.
- Verify the certificate and key generation steps. Ensure that the IP Subject Alternative Name (SAN) is specified correctly.

Certificate is not valid

If the Cloud-Delivered Firewall Management Center displays the error "Request was made for (IP), but the certificate is not valid for (IP)" then verify the server certificate and key generation steps.

- Ensure that the IP SAN is correctly specified in the server certificate. If the configuration applies to more than one Firewall Threat Defense device, you must specify all the devices in the IP SAN field.
- If you are using dial-out mode, ensure that the client IP is specified in the server certificate.

Failed to generate response object

If you receive "Failed to generate response object, did not receive any data" error, the gNMI input plug-in is waiting for metric export. This is the sample response that appears when the telegraf is restarting:

```
root@cronserver:/home/secanup/openconfig-test# gnmic -a $ADDRESS:$PORT --tls-cert $CLIENTCERT
 --tls-ca $CACERT --tls-key $CLIENTKEY -u $USER -p $PASS sub --mode once --path
"openconfig-system/system/memory"
rpc error: code = Aborted desc = Error in gnmi_server: failed to generate response object.did
 not receive any data
Error: one or more requests failed
```

Wait for the gNMI input plug-in to restart and retry your request.

Restart telegraf

When telegraf is not responding, restart the process using this command on the Firewall Threat Defense CLI console:

```
pmtool restartbyid hmdaemon
```

Get current status of gNMI server

When OpenConfig streaming telemetry is enabled, to know the status of the gNMI server, run this command using the Firewall Threat Defense CLI console:

```
curl localhost:9275/OpenConfig/status
```

This is the sample response to the command:

```
root@firepower:/home/admin# curl localhost:9275/openconfig/status
Mode (Dialin/Dialout): DialIn
Subscription Details:
    Active Subscription Details:
        Stream Mode Subscription Details:
            Total Stream Subscription Request Count: 1
            'Ip of Collector- Subscribe paths:'
                172.16.0.101:45826:
                - /openconfig-system/system/state/hostname
            Sample Subscription Count: 1
            On Change Subscription Count: 0
        Once Mode Subscription Details:
            Total Subscription Request Count: 0
            Total Subscription Count: 0
            'Ip of Collector- Subscribe paths:': {}
    Total Subscription Details:
        Stream Mode Subscription Details:
            Total Stream Subscription Request Count: 1
            'Ip of Collector- Subscribe paths:':
                 172.16.0.101:45826:
                 - /openconfig-system/system/state/hostname
        Sample Subscription Count: 1
        On Change Subscription Count: 0
    Once Mode Subscription Details:
        Total Subscription Request Count: 0
        Total Subscription Count: 0
        'Ip of Collector- Subscribe paths:': {}
```

system auditing

The appliances that are part of the system generate an audit record for each user interaction with the web interface.

Audit records

Cloud-Delivered Firewall Management Center logs read-only auditing information for user activity. Audit logs appear in a standard event view that allows you to view, sort, and filter audit log messages based on any item in the audit view. You can easily delete audit information, generate reports, and review detailed records of the changes that users make.

Audit record storage and limitations

The audit log stores a maximum of 100,000 entries. When the number of audit log entries exceeds 100,000, the appliance prunes the oldest records from the database to reduce the number to 100,000.

The audit logs do not display the user or the source IP for login errors:

- When wrong password is used, the source IP is not displayed.
- When the user account does not exist, both source IP and the user are not displayed.
- If the attempt for an LDAP user fails, no audit log is triggered.

Related Topics

Guidelines for SSO configuration on Cloud-Delivered Firewall Management Center

Audit log workflow fields

This table describes the audit log fields that can be viewed and searched.

Table 25: Audit log fields

Field	Description
Time	Time and date that the appliance generated the audit record.
User	User name of the user that triggered the audit event.
Subsystem	In a few cases where a menu path is not relevant, the Subsystem field displays only the event type. For example, **Login** classifies user login attempts.
Message	The action the user performed or the button the user clicked on the page. For example, `Page View` signifies that the user simply viewed the page indicated in the Subsystem, while `Save` means that the user clicked the **Save** button on the page. Changes made to the system appear with a **Compare icon** that you can click to see a summary of the changes.

Field	Description
Source IP	IP address associated with the host used by the user. Note: When searching this field you must type a specific IP address; you cannot use IP ranges when searching audit logs.
Domain	The current domain of the user when the audit event was triggered. This field is only present if you have ever configured the Cloud-Delivered Firewall Management Center for multitenancy.
Configuration Change (search only)	Specifies whether to view audit records of configuration changes in the search results. (`yes` or `no`)
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

The audit events table view

You can customize the layout of the audit events table view or filter events by specific field value.

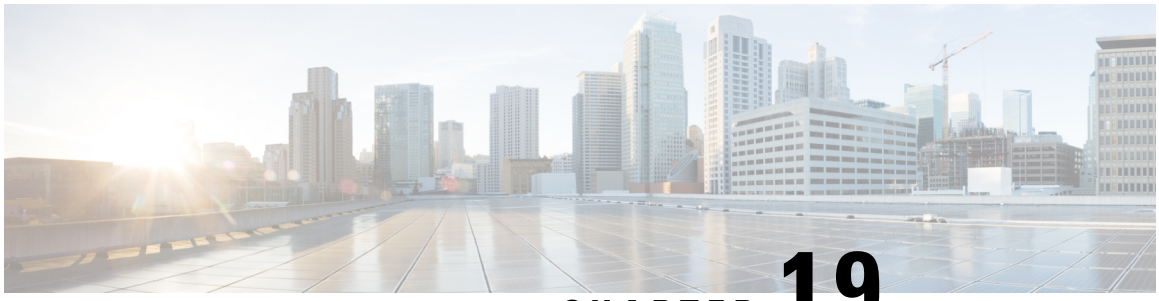
Use these options to customize the audit events table view:

- You can change the layout of the event view, or constrain the events in the view by a field value.
- To disable a column, click **Close** (✕) in the desired column heading. Click **Apply** in the pop-up window. When you disable a column, it is disabled for the duration of your session (unless you add it back later).

Note When you disable the first column, the Count column is added.

- To hide or show other columns, or to add a disabled column back to the view, select or clear the appropriate check boxes before you click **Apply**.
- Clicking a value within a row in a table view constrains the table view and does not drill down to the next page in the workflow.

Tip Table views always include "Table View" in the page name.



CHAPTER 19

Troubleshooting

The following topics describe ways to diagnose problems you may encounter:

- Guidelines and best practices for troubleshooting, on page 419
- System messages, on page 420
- View basic system information, on page 422
- Manage system messages, on page 423
- Memory usage thresholds for health monitor alerts, on page 427
- Disk usage and drain of events health monitor alerts, on page 428
- Clear disk space, on page 432
- Health monitor reports for troubleshooting, on page 433
- Enhanced troubleshooting experience using Cisco RADKit integration, on page 435
- General troubleshooting, on page 439
- Connection-based troubleshooting, on page 439
- Advanced troubleshooting for the Secure Firewall Threat Defense device, on page 440
- Feature-specific troubleshooting, on page 453

Guidelines and best practices for troubleshooting

Follow these best practices to troubleshoot system issues effectively and prepare for potential support escalation:

- Before you make changes to try to fix a problem, generate a troubleshooting file to capture the original problem. See Health monitor reports for troubleshooting, on page 433 and its subsections.

  You may need this troubleshooting file if you need to contact Cisco TAC for support.

- Start your investigation by looking at error and warning messages in the Message Center. See System messages, on page 420

- Look for applicable Tech Notes and other troubleshooting resources under the "Troubleshoot and Alerts" heading on the product documentation page for your product.

- During the troubleshooting process, as several commands are executed simultaneously, the CPU usage becomes high. We recommend that you perform troubleshooting during periods of lower network traffic and fewer users.

- The time taken to troubleshoot the system depends on the configuration settings and the number of devices configured. This duration can range from nine to sixty minutes.

System messages

When you need to track down problems occurring in the system, the Message Center is the place to start your investigation. This feature allows you to view the messages that the system continually generates about system activities and status.

Message center access and display

To open the Message Center, click the System Status icon, next to the **Deploy** button on the home page. This icon changes based on the system status:

- **Error** ( ): Indicates one or more errors and any number of warnings are present on the system.
- **Warning** ( ): Indicates one or more warnings and no errors are present on the system.
- **Success** ( ): Indicates no warnings or errors are present on the system.

If a number is displayed with the icon, it indicates the total current number of error or warning messages.

To close the Message Center, click anywhere outside of it within the web interface.

The web interface also displays pop-up notifications in response to your actions and ongoing system activities. Some pop-up notifications automatically disappear after five seconds, while others are "sticky," meaning they display until you explicitly dismiss them by clicking **Dismiss** ( ). Click the **Dismiss** link at the top of the notifications list to dismiss all notifications at once.

Tip Hovering your cursor over a non-sticky pop-up notification causes it to be sticky.

The system determines which messages it displays to users in pop-up notifications and the Message Center based on their licenses, domains, and access roles.

Message types

The Message Center displays messages reporting system activities and status organized into three different tabs.

The Message Center organizes system status messages into these tabs:

- **Deployments**: Displays current status related to configuration deployment for each appliance in your system, grouped by domain.
- **Upgrades**: Displays the current status related to software upgrade tasks for the managed devices.
- **Health**: Displays current health status information for each appliance in your system, grouped by domain.
- **Tasks**: Displays the status of long-running tasks such as configuration backups, update installation, etc,.

Deployment status

The system reports these deployment statuses. You can get additional detail about the deployment jobs by clicking **Show History**.

- Running (**Spinning**): The configuration is in the process of deploying.
- **Success**: The configuration has successfully been deployed.
- **Warning** ( ) : Warning deployment statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Failure**: The configuration has failed to deploy. Failed deployments contribute to the message count displayed with the **Error System Status icon**.

Upgrade status

The system reports these upgrade status values on the Upgrade tab:

- **In progress**: Indicates that the upgrade task is in progress.
- **Completed**: Indicates that the software upgrade task is completed successfully.
- **Failed**: Indicates that the software upgrade task has failed to complete.

Health status

This tab displays current health status information for each appliance in your system, grouped by domain. Health status is generated by health modules as described in Health monitoring in Cloud-Delivered Firewall Management Center, on page 371. The system reports the following health status values on this tab:

- **Warning** ( ): Indicates that warning limits have been exceeded for a health module on an appliance and the problem has not been corrected. The Health Monitoring page indicates these conditions with a **Yellow Triangle** ( ). Warning statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Critical** ( ): Indicates that critical limits have been exceeded for a health module on an appliance and the problem has not been corrected. The Health Monitoring page indicates these conditions with a **Critical** ( ) icon. Critical statuses contribute to the message count displayed with the **Error System Status icon**.
- **Error** ( ): Indicates that a health monitoring module has failed on an appliance and has not been successfully re-run since the failure occurred. The Health Monitoring page indicates these conditions with an **Error icon**. Error statuses contribute to the message count displayed with the **Error System Status icon**.

You can click the links in the Health tab to view related detailed information on the Health Monitoring page. If there are no current health status conditions, the Health tab displays no messages.

Task status

Certain tasks (such as configuration backups or update installation) can require some time to complete. This tab displays the status of these long-running tasks, and can include tasks initiated by you or, if you have

appropriate access, other users of the system. The tab presents messages in reverse chronological order based on the most recent update time for each message. Some task status messages include links to more detailed information about the task in question. The system reports the following task status values on this tab:

- **Waiting**(): Indicates a task that is waiting to run until another in-progress task is complete. This message type displays an updating progress bar.
- **Running**: Indicates a task that is in-progress. This message type displays an updating progress bar.
- **Retrying**: Indicates a task that is automatically retrying. Note that not all tasks are permitted to try again. This message type displays an updating progress bar.
- **Success**: Indicates a task that has completed successfully.
- **Failure**: Indicates a task that did not complete successfully. Failed tasks contribute to the message count displayed with the **Error System Status icon**.
- **Stopped or Suspended**: Indicates a task that was interrupted due to a system update. Stopped tasks cannot be resumed. After normal operations are restored, start the task again.
- Skipped: A process in progress prevented the task from starting. Try again to start the task.

New messages appear in this tab as new tasks are started. As tasks complete (status success, failure, or stopped), this tab continues to display messages with final status indicated until you remove them. We recommend you remove messages to reduce clutter in the Tasks tab as well as the message database.

Message management

From the Message Center you can perform these tasks:

- Choose to display pop-up notifications.
- Display more task status messages from the system database (if any are available that have not been removed).
- Download a report of all the task manager notifications.
- Remove individual task status messages. (This affects all users who can view the removed messages.)
- Remove task status messages in bulk. (This affects all users who can view the removed messages.)

Tip Cisco recommends that you periodically remove accumulated task status messages from the Task tab to reduce clutter in both the display and the database. When the number of messages in the database approaches 100,000, the system automatically deletes task status messages that you have removed.

View basic system information

The About page displays information about your appliance, including the model, serial number, and version information for various components of the system. It also includes Cisco copyright information.

Procedure

Step 1 Click **Help** (?) in the toolbar at the top of the page.

Step 2 Choose **About**.

View appliance information

Procedure

Choose **System** (⚙) > **Configuration**.

Manage system messages

Use the Message Center to monitor and manage different types of system notifications including configuration deployments, device upgrades, system health alerts, and long-running tasks.

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Choose from the available notification types:

- Click **Deployments** to view messages related to configuration deployments. See View deployment messages, on page 424. You must be an Admin user or have the **Deploy Configuration to Devices** permission to view these messages.
- Click **Upgrades** to view messages related to device upgrade tasks. See Viewing Upgrade Messages. See Viewing Upgrade Messages. You must be an Admin user or have **Updates** permission to view these messages.
- Click **Health** to view messages related to the health of your Cloud-Delivered Firewall Management Center and the devices registered to it. See View health messages, on page 425. You must be an Admin user or have the **Health** permission to view these messages.

  You can navigate to the Health Monitor page by clicking the **Health monitor** link.

- Click **Tasks** to view or manage messages related to long-running tasks. See View task messages, on page 426 or Manage task messages, on page 427. Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission. You can remove the completed tasks from the notification by clicking the **Remove completed tasks** link.
- Click the **Download Report** icon to generate a report of all the notifications in the task manager. Choose **Download CSV** or **Download PDF** to download the report.

- Click **Show Notifications** slider to enable or disable pop-up notification display.

View deployment messages

View deployment messages to monitor the status of configuration deployments and access detailed deployment history for troubleshooting.

If you are an Admin user or have the **Deploy Configuration to Devices** permission, you can view deployment messages and history.

Complete these steps to view deployment messages and deployment history:

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Deployments**.

Step 3 You have the following choices:

- Click **total** to view all current deployment statuses.
- Click a status value to view only messages with that deployment status.
- Hover your cursor over the time elapsed indicator for a message (for example, **1m 5s**) to view the elapsed time, and start and stop times for the deployment.

Step 4 Click **show deployment history** to view more detailed information about the deployment jobs.

The Deployment History table lists the deployment jobs in the left column in reverse chronological order.

a) Select a deployment job.

The table in the right column shows each device that was included in the job, and the deployment status per device.

b) To view responses from the device, and commands sent to the device during deployment, click download in the **Transcript** column for the device.

The transcript includes the following sections:

- **Snort Apply**: If there are any failures or responses from Snort-related policies, messages appear in this section. Normally, this section is empty.
- **CLI Apply**: This section covers features that are configured using commands sent to the Lina process.
- **Infrastructure Messages**: This section shows the status of different deployment modules.

In the **CLI Apply** section, the deployment transcript lists commands sent to the device as well as any responses returned from the device. These responses can be informative messages or error messages. If a deployment fails, review the error messages to understand which commands were problematic. If you use FlexConfig policies to configure features, these messages can help you identify and correct issues in your FlexConfig scripts.

Note

The transcript does not distinguish between commands sent for managed features and those generated from FlexConfig policies.

For example, this sequence shows that the Cloud-Delivered Firewall Management Center sent commands to configure GigabitEthernet0/0 with the logical name outside. The device responded that it automatically set the security level to 0. The Firewall Threat Defense does not use the security level for anything.

```
========= CLI APPLY =========

FMC >> interface GigabitEthernet0/0
FMC >>  nameif outside
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

You can view deployment messages, monitor deployment statuses, and access detailed deployment history, including device transcripts to troubleshoot deployment issues.

View upgrade messages

View and monitor upgrade messages to track the status of current upgrade tasks and access detailed information about upgrade processes.

Only Admin users or users with the **Updates** permission can view these messages.

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Upgrades**.

Step 3 You can do the following:

- Click **total** to view all current upgrade tasks.
- Click on a status value to see messages with only that status.
- Click **Device Management** for more details on the upgrade task.

View health messages

Monitor system status, identify issues, and access detailed health information for troubleshooting and maintenance.

You must be an Admin user or have the **Health** permission to view health messages. Using the Message Center provides a streamlined way to monitor the system's health and filter messages by severity.

Complete these steps to view health messages in the Message Center:

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Health**.

Step 3 You have the following choices:

- Click **total** to view all current health statuses. You also see the severity levels: warning, critical, and error.
- Click on a status value to view only messages with that status.
- Hover your cursor over the relative time indicator for a message (for example, **3 day(s) ago**) to view the time of the most recent update for that message.
- To view detailed health status information for a particular message, click the message.
- To view complete health status on the Health Monitoring page, click **Health Monitor**.

You can now view, filter, and access detailed health status messages for monitoring and troubleshooting your system.

View task messages

View task status messages to monitor the progress and completion status of various system tasks.

You can view your tasks. But, to see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Tasks**.

Step 3 You have these choices:

- Click **total** to view all current task statuses. To view tasks based on a specific status—such as waiting, running, retrying, success, or failure—click its corresponding label.
- Click a status value to view only messages for tasks with that status.

  Note
  Messages for stopped tasks appear only in the total list of task status messages. You cannot filter on stopped tasks.

- Hover your cursor over the relative time indicator for a message (e.g., **3 day(s) ago**) to view the time of the most recent update for that message.
- Click any link within a message to view more information about the task.
- If more task status messages are available for display, click **Fetch more messages** at the bottom of the message list to retrieve them.

Manage task messages

You can view your tasks. But, to see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click **Tasks**.

Step 3 Select one of the available management options:

- If more task status messages are available for display, click **Fetch more messages** at the bottom of the message list to retrieve them.
- To remove a single message for a completed task (status stopped, success, or failure), click **Remove** ( ) next to the message.
- To remove all messages for all tasks that have completed (status stopped, success, or failure), filter the messages on **total** and click **Remove all completed tasks**.
- To remove all messages for all tasks that have completed successfully, filter the messages on **success**, and click **Remove all successful tasks**.
- To remove all messages for all tasks that have failed, filter the messages on **failure**, and click **Remove all failed tasks**.

Memory usage thresholds for health monitor alerts

The memory usage health module compares memory usage on an appliance to the limits configured for the module and alerts when usage exceeds the levels. The module monitors data from managed devices and from the Cloud-Delivered Firewall Management Center itself.

Two configurable thresholds for memory usage, critical and warning, can be set as a percentage of memory used. When these thresholds are exceeded, a health alarm is generated with the severity level specified. However, the health alarm system does not calculate these thresholds in an exact manner.

With high memory devices, certain processes are expected to use a larger percentage of total system memory than in a low memory footprint device. The design is to use as much of the physical memory as possible while leaving a small value of memory free for ancillary processes.

Compare two devices, one with 32 GB of memory and one with 4 GB of memory. In the device with 32 GB of memory, 5% of memory (1.6GB) is a much larger value of memory to leave for ancillary processes than in the device with 4 GB of memory (5% of 4GB = 200MB).

To account for the higher percentage use of system memory by certain processes, the Cloud-Delivered Firewall Management Center calculates the total memory to include both total physical memory and total swap memory. Thus the enforced memory threshold for the user-configured threshold input can result in a Health Event where the **Value** column of the event does not match the value that was entered to determine the exceeded threshold.

Beginning with version 7.4.1, the memory usage health module calculates memory usage by taking into account the free memory available, swap memory available, and buffer cache. To avoid premature memory usage health alerts, it is recommended not to exceed the warning and critical alarm thresholds of 88% and 90%, respectively.

The following table shows examples of user-input thresholds and the enforced thresholds, depending on the installed system memory.



Note The values in this table are examples. You can use this information to extrapolate thresholds for devices that do not match the installed RAM shown here, or you can contact Cisco TAC for more precise threshold calculations.

Table 26: Memory usage thresholds based on installed RAM

User-Input Threshold Value	Enforced Threshold Per Installed Memory (RAM)			
	4 GB	6 GB	32 GB	48 GB
10%	10%	34%	72%	81%
20%	20%	41%	75%	83%
30%	30%	48%	78%	85%
40%	40%	56%	81%	88%
50%	50%	63%	84%	90%
60%	60%	70%	88%	92%
70%	70%	78%	91%	94%
80%	80%	85%	94%	96%
90%	90%	93%	97%	98%
100%	100%	100%	100%	100%

Disk usage and drain of events health monitor alerts

This topic describes the symptoms and troubleshooting guidelines for two health alerts generated by the Disk Usage health module:

- Frequent Drain of Events
- Drain of Unprocessed Events

Disk usage overview

The disk usage health module compares disk usage on a managed device's hard drive and malware storage pack to the limits configured for the module and alerts when usage exceeds the percentages configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds.

The disk manager process manages the disk usage of a device. Each type of file monitored by the disk manager is assigned with a silo. Based on the amount of disk space available on the system the disk manager computes a High Water Mark (HWM) and a Low Water Mark (LWM) for each silo.

To display detailed disk usage information for each part of the system, including silos, LWMs, and HWMs, use the **show disk-manager** command.

Examples

The following is an example of the disk manager information:

```
> show disk-manager
Silo                                Used        Minimum     Maximum
Temporary Files                     0 KB        499.197 MB  1.950 GB
Action Queue Results                0 KB        499.197 MB  1.950 GB
User Identity Events                0 KB        499.197 MB  1.950 GB
UI Caches                           4 KB        1.462 GB    2.925 GB
Backups                             0 KB        3.900 GB    9.750 GB
Updates                             0 KB        5.850 GB    14.625 GB
Other Detection Engine              0 KB        2.925 GB    5.850 GB
Performance Statistics              33 KB       998.395 MB  11.700 GB
Other Events                        0 KB        1.950 GB    3.900 GB
IP Reputation & URL Filtering       0 KB        2.437 GB    4.875 GB
Archives & Cores & File Logs        0 KB        3.900 GB    19.500 GB
Unified Low Priority Events         1.329 MB    4.875 GB    24.375 GB
RNA Events                          0 KB        3.900 GB    15.600 GB
File Capture                        0 KB        9.750 GB    19.500 GB
Unified High Priority Events        0 KB        14.625 GB   34.125 GB
IPS Events                          0 KB        11.700 GB   29.250 GB
```

Health alert format

When the Health Monitor process on the Cloud-Delivered Firewall Management Center runs (once every 5 minutes or when a manual run is triggered), the Disk Usage module looks into the diskmanager.log file and, if the correct conditions are met, the health alert is triggered.

The structures of these health alerts are as follows:

- Frequent drain of *SILO NAME*
- Drain of unprocessed events from *SILO NAME*

For example,

- Frequent drain of Low Priority Events
- Drain of unprocessed events from Low Priority Events

Its possible for any silo to generate a *Frequent drain of SILO NAME* health alert. However, the most commonly seen are the alerts related to events. Among the event silos, the *Low Priority Events* are often seen because device generates this type of events frequently.

A *Frequent drain of <SILO NAME>* event has a **Warning** severity level when seen in relation to an event-related silo, because events will be queued to be sent to the Cloud-Delivered Firewall Management Center. For a non-event related silo, such as the *Backups* silo, the alert has a **Critical** severity level because this information is lost.

Note Only event silos generate a *Drain of unprocessed events from SILO NAME* health alert. This alert always has a **Critical** severity level.

Additional symptoms besides the alerts can include:

- Slowness on the Cloud-Delivered Firewall Management Center user interface
- Loss of events

Common troubleshoot scenarios

A *Frequent drain of SILO NAME* event is caused by too much input into the silo for its size. In this case, the disk manager drains (purges) that file at least twice in the last 5-minute interval. In an event type silo, this is typically caused by excessive logging of that event type.

A *Drain of unprocessed events of SILO NAME* health alert is caused by a bottleneck in the event processing path.

There are three potential bottlenecks with respect to these Disk Usage alerts:

- Excessive logging — The EventHandler process on Firewall Threat Defense is oversubscribed (it reads slower than what Snort writes).
- Sftunnel bottleneck — The Eventing interface is unstable or oversubscribed.
- SFDataCorrelator bottleneck — The data transmission channel between the Cloud-Delivered Firewall Management Center and the managed device is oversubscribed.

Excessive logging

One of the most common causes for the health alerts of this type is excessive input. The difference between the Low Water Mark (LWM) and High Water Mark (HWM) gathered from the **show disk-manager** command shows how much space there is available to take on that silo to go from LWM (freshly drained) to the HWM value. If there are frequent drain of events (with or without unprocessed events), review the logging configuration.

- Check for double logging — Double logging scenarios can be identified if you look at the correlator *perfstats* on the Cloud-Delivered Firewall Management Center:

  ```
  admin@FMC:~$ sudo perfstats -Cq < /var/sf/rna/correlator-stats/now
  ```

- Check logging settings for the ACP — Review the logging settings of the Access Control Policy (ACP). If the logging setting includes both "Beginning" and "End" of connection, modify the setting to log only the end to reduce the number of events.

Communications Bottleneck – Sftunnel

Sftunnel is responsible for encrypted communications between the Cloud-Delivered Firewall Management Center and the managed device. Events are sent over the tunnel to the Cloud-Delivered Firewall Management Center. Connectivity issues and/or instability in the communication channel (sftunnel) between the managed device and the Cloud-Delivered Firewall Management Center can be due to:

- Sftunnel is down or is unstable (flaps).

  Ensure that the Cloud-Delivered Firewall Management Center and the managed device have reachability between their management interfaces on TCP port 8305.

  The sftunnel process should be stable and should not restart unexpectedly. Verify this by checking the `/var/log/message` file and search for messages that contain the **sftunneld** string.

- Sftunnel is oversubscribed.

  Review trend data from the Heath Monitor and look for signs of oversubscription of the Cloud-Delivered Firewall Management Center's management interface, which can be a spike in management traffic or a constant oversubscription.

  Use as a secondary management interface for eventing. To use this interface, you must configure its IP address and other parameters at the Firewall Threat Defense CLI using the **configure network management-interface** command.

Communications Bottleneck – SFDataCorrelator

The SFDataCorrelator manages data transmission between the Cloud-Delivered Firewall Management Center and the managed device; on the Cloud-Delivered Firewall Management Center, it analyzes binary files created by the system to generate events, connection data, and network maps. The first step is to review the **diskmanager.log** file for important information to be gathered, such as:

- The frequency of the drain.
- The number of files with Unprocessed Events drained.
- The occurrence of the drain with Unprocessed Events.

Each time the disk manager process runs it generates an entry for each of the different silos on its own log file, which is located under **[/ngfw]/var/log/diskmanager.log**. Information gathered from the diskmanager.log (in CSV format) can be used to help narrow the search for a cause.

Additional troubleshooting steps:

- The command **stats_unified.pl** can help you to determine if the managed device does have some data which must be sent to Cloud-Delivered Firewall Management Center. This condition can happen when the managed device and the Cloud-Delivered Firewall Management Center experience a connectivity issue. The managed device stores the log data on to a hard drive.

  ```
  admin@FMC:~$ sudo stats_unified.pl
  ```

- The **manage_proc.pl** command can reconfigure the correlator on the Cloud-Delivered Firewall Management Center side.

  ```
  root@FMC:~# manage_procs.pl
  ```

Before you contact Cisco TAC

It is highly recommended to collect these items before you contact Cisco TAC:

- Screenshots of the health alert seen.
- Troubleshoot file generated from the Cloud-Delivered Firewall Management Center.
- Troubleshoot file generated from the affected managed device.
- Date and Time when the problem was first seen.
- Information about any recent changes done to the policies (if applicable).
- The output of the stats_unified.pl command as described in Communications Bottleneck — SFDataCorrelator, on page 431.

Clear disk space

Low disk space can reduce the performance of the managed Firewall Threat Defense devices. It can also prevent device upgrades and increase the risk of accidentally deleting important files when trying to recover space.

You can free up disk space by removing the following temporary files from your Firewall Threat Defense device:

- **Backup files**: These are the Firewall Threat Defense device backup files. Device backups include only the configuration data. For more information, refer to Back Up Managed Devices, on page 462.
- **Content updates**: This includes SRU, VDB, and GeoDB update files. For more information, refer to About System Updates.
- **Troubleshooting files**: These are log files generated for troubleshooting purposes. For more information, refer to Health Monitor Reports for Troubleshooting.

When you clear the disk space, the system retains the latest file and deletes all older files.

Caution Clearing disk space will permanently delete the selected files.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 In the left navigation pane, click the device for which you want to clear disk space.

Step 3 In the **Disk Usage** widget, click **Clear disk space**.

Step 4 Choose the type of temporary files that you want to remove by checking the corresponding check box.

Step 5 Click **Clear disk space**.

Step 6 Click **Clear** to continue.

Check the message center for the progress of the disk cleanup task. When the task is complete, the **Disk Usage** widget will display the updated storage data.

Health monitor reports for troubleshooting

In some cases, if you have a problem with your appliance, support may ask you to supply troubleshooting files to help them diagnose the problem. The system can produce troubleshooting files with information targeted to specific functional areas, as well as advanced troubleshooting files you retrieve in cooperation with Support. You can select any of the options listed in the table below to customize the contents of a troubleshooting file for a specific function.



Note Some options overlap in terms of the data they report, but the troubleshooting files will not contain redundant copies, regardless of what options you select.

Table 27: Selectable troubleshoot options

This option...	Reports...
Snort Performance and Configuration	data and configuration settings related to Snort on the appliance
Hardware Performance and Logs	data and logs related to the performance of the appliance hardware
System Configuration, Policy, and Logs	configuration settings, data, and logs related to the current system configuration of the appliance
Detection Configuration, Policy, and Logs	configuration settings, data, and logs related to detection on the appliance
Interface and Network Related Data	configuration settings, data, and logs related to inline sets and network configuration of the appliance
Discovery, Awareness, VDB Data, and Logs	configuration settings, data, and logs related to the current discovery and awareness configuration on the appliance
Upgrade Data and Logs	data and logs related to prior upgrades of the appliance
All Database Data	all database-related data that is included in a troubleshoot report
All Log Data	all logs collected by the appliance database
Network Map Information	current network topology data

Generate troubleshooting files for specific system functions

You can generate and download customized troubleshooting files that you can send to support.

Before you begin

You must be an Admin, Maintenance, Security Analyst, or Security Analyst (Read Only) user to perform this task.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**, click the device in the left panel, then **View System & Troubleshoot Details**, and then click **Generate Troubleshooting Files**.

Note

- Troubleshooting files generated from the Cloud-Delivered Firewall Management Center web interface are stored on the Cloud-Delivered Firewall Management Center. Only the latest troubleshooting file for each appliance is stored.
- Troubleshooting files generated from the CLI are stored locally and are not overwritten.
- If you initiate the troubleshooting log generation while one is already in the queue, the new request is dropped.

Step 2 Choose **All Data** to generate all possible troubleshooting data, or check individual boxes as described in View task messages, on page 426.

Step 3 Click **Generate**.

Step 4 View task messages in the Message Center; see View task messages, on page 426.

Step 5 Find the task that corresponds to the troubleshooting files you generated.

Step 6 After the appliance generated the troubleshooting files and the task status changes to `Completed`, click **Click to retrieve generated files**.

Step 7 Follow your browser's prompts to download the file. (The troubleshooting files are downloaded in a single `.tar.gz` file.)

Step 8 Follow the directions from Support to send the troubleshooting files to Cisco.

Download advanced troubleshooting files

Download troubleshooting files and share them with Cisco support for advanced troubleshooting assistance when requested.

Before you begin

You must have **Admin**, **Maintenance**, **Security Analyst**, or **Security Analyst (Read Only)** user roles to perform this task.

Procedure

Step 1 View the health monitor for the appliance.

Step 2 Choose **System** (⚙) > **Health** > **Monitor**.

Step 3 Click a device in the left pane for which you dowload the troubleshooting files.

Step 4 Click **View System & Troubleshoot Details**, and then click **Advanced Troubleshooting.**

Step 5 In **File Download** field, enter the name of the file to download. Cisco support provides this file name.

Step 6 Click **Download**.

Step 7 Follow your browser's prompts to download the file.

Note
For managed devices, the system renames the file by prepending the device name to the file name.

Step 8 Follow the directions from Cisco support to send the troubleshooting files to Cisco.

Enhanced troubleshooting experience using Cisco RADKit integration

The Cisco Remote Automation Development Kit (RADKit) is a network-wide orchestrator designed to provide Cisco support engineers with secure access to network devices for diagnosing and troubleshooting system issues. You can install RADKit on a remote server and integrate its services with Cloud-Delivered Firewall Management Centers and their managed Firewall Threat Defense devices such that the data required for diagnosing issues are available automatically.

RADKit integration capabilities

Note RADKit supports only the instances of a cluster device, not the chassis itself.

Integrating RADKit with the Cloud-Delivered Firewall Management Center enables you to:

- Access the Cloud-Delivered Firewall Management Center and Firewall Threat Defense device remotely.
- Assign an engineer with controlled access to the Cloud-Delivered Firewall Management Center and Firewall Threat Defense devices for diagnosing configuration discrepancies and hardware issues.
- Leverage automation capabilities to collect diagnostic data and analyze problems.
- Navigate to the previous sessions, view, and download all the operations logs as an archive file.

Guidelines and limitations for RADKit integration

Follow these guidelines when configuring RADKit integration to ensure proper functionality:

- If the Cloud-Delivered Firewall Management Center is restored with a backup file of another Cloud-Delivered Firewall Management Center, the devices mapped to the current authorizations are cleared. You must edit the authorization and repeat device selection in the **Authorization** page.
- If you are using HTTP proxy, configure the proxy in the Cloud-Delivered Firewall Management Center before enabling the RADKit service for a sucessful enrolment of the service. In case of incorrect proxy configuration at the time of RADKit enablement, disable the RADKit service, after configuring the correct proxy, re-enable for sucessful enrolment of the service.

Enroll RADKit service

You can enroll the RADKit service with SSO, which establishes connectivity with the RADKit client.

Before you begin

If you intend to use HTTP proxy, configure the proxy in the management center before enabling the RADKit service for sucessful enrolment of the service.

Procedure

Step 1 Choose **Devices** > **Troubleshoot** > **Remote Diagnostics**.

Step 2 To enroll the service in the RADKit cloud, click the **Enable the RADKit service** toggle button, and then click **Enroll with SSO** .

Step 3 Enter your email address (with domain name), and click **Submit** .

Note
You cannot enroll multiple times within a short period using the same email ID. Each time that you enroll, a new certificate is generated for the email ID account. Creating multiple certificates for the same user account is not efficient. Hence, RADKit Cloud restricts the number of issued certificates per user to 12 new Service IDs for a six-month duration.

Step 4 Complete the SSO process.

Step 5 In the **Cisco RADKit Access** page, click **Accept** .

After enrollment, the page displays the success message and the token ID.

Step 6 (Optional) To disable the service, click the **Disable the RADKit service** toggle button.

Caution
When you disable the service, all the associated data and existing authorizations will be lost.

Manage RADKit service authorization

Before you begin

Enroll to RADKit service. For information about the procedure to enroll, see Enroll RADKit service, on page 436.

You can manage authorizations by enabling the TAC engineers to access the devices remotely for troubleshooting. You can also grant access to remote users to either access all the devices' inventory or only of selected devices for a specific duration.

Note The authorizations can be added and viewed only from the Global domain.

Procedure

Step 1 Choose **Devices** > **Troubleshoot** > **Remote Diagnostics**.

The **Remote Diagnostics** page lists the current authorizations. Ensure that you have enrolled to the RADKit service.

Step 2 To create a new authorization, click **Create New Authorization** .

Step 3 Under **Setup Authorization** , enter the support engineer's email address. The Email address must not exceed 40 characters.

Note

- The support engineer's email address is unique to an authorization record. Hence, you cannot modify this value after creating authorization.
- You cannot create an authorization account for the same email address.

Step 4 Click **Next** .

Step 5 Under **Device Selection**, specify the devices and the Cloud-Delivered Firewall Management Center that the support engineer can access. Select the relevant option:

- **Grant access to all devices**: Click this radio button, and then check the **I agree to provide access to all devices** check box to agree for granting access to all devices and the Cloud-Delivered Firewall Management Center.
- **Grant access to specific devices**: Click this radio button, and then from the **Select Devices** drop-down list, choose the device IP addresses and the Cloud-Delivered Firewall Management Center that the support engineer can access. Ensure to check the **I agree to provide access to the selected devices** check box.

  Note
  Only devices running Version 7.7 and later are supported and are available for selection.

Step 6 Click **Next**.

Step 7 Under **Schedule the Access**, specify the timeline during which the support engineer can access the devices:

- **Now** : Click this radio button, and then under **Duration** fields, specify the time duration for granting access. To set the duration preset, click the **Time presets** links, namely, 1 hour, 6 hours, 12 hours, 1 day, or 1 week.
- **Grant access until revoke** : Click this radio button to grant perpetual access to the devices.

  Note
  To revoke access, in the **Current Authorizations** grid, click the respective revoke icon under the **Actions** column.

Step 8 Under **Summary** , preview the selections you have made for the authorization account. To modify the choices made, click the Edit icon and make the corresponding changes.

Note
The support engineer's email address is unique to an authorization record. Hence, you cannot modify the email address.

Step 9 Click **Create** .

In the subsequent window, a confirmation message for the new authorization with the user email address and the RADKit service id used in the RADKit client is displayed.

Step 10 (Optional) To create another authorization, click **Create Another Authorization** .

Step 11 To exit, click **Close** .

What to do next

The **Remote Diagnostics** window displays a list of current authorizations. You can perform the following actions:

- Edit the settings.
- Revoke the authorization.
- Enable sudo access for the devices and Cloud-Delivered Firewall Management Center.

Enable sudo access for devices

When you provide sudo device access, regular users can have administrative rights to the devices. Irrespective of the user role, when you enable sudo access for a device, all the users of the device can enjoy this privilege.

Before you begin

Ensure that you have added the required devices in the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Troubleshoot** > **Remote Diagnostics**.

Step 2 To enable sudo access to devices, click the **Device Sudo Access** tab.

The devices deployed in the Cloud-Delivered Firewall Management Center and the Cloud-Delivered Firewall Management Center are listed.

Step 3 To enable sudo access for a device, click the respective toggle button under **Sudo Status**.

Step 4 For bulk action, check the check boxes adjacent to the devices, and click **Enable**.

Step 5 To disable sudo access, click the respective toggle button under **Sudo Status**.

Download session logs

The **Previous Sessions** tab displays a list of log files for the operations performed by the RADKit client in the Firewall Threat Defense devices and the Cloud-Delivered Firewall Management Center. These session logs can be downloaded as an archive. The RADKit sessions' logs are added to the Cloud-Delivered Firewall Management Center log rotate mechanism and fall under the default log rotate mechanism as that of other Cloud-Delivered Firewall Management Center log files.

Procedure

Step 1 Choose **Devices** > **Troubleshoot** > **Remote Diagnostics** and click the **Previous Sessions** tab.

Step 2 Search for a specific session log by entering the search string in the **Past Session Logs** search bar.

Step 3 Download the logs as an archive using **Download All Logs**.

Note
You can download the logs from the **Current Authorization** page too.

General troubleshooting

An internal power failure (hardware failure, power surge, and so on) or an external power failure (unplugged cord) can result in an ungraceful shutdown or reboot of the system. This can result in data corruption.

Connection-based troubleshooting

Connection-based troubleshooting or debugging provides uniform debugging across modules to collect appropriate logs for a specific connection. It also supports level-based debugging up to seven levels and enables uniform log collection mechanism across modules. Connection-based debugging supports the following:

- A common connection-based debugging subsystem to troubleshoot issues in Firewall Threat Defense
- Uniform format for debug messages across modules
- Persistent debug messages across reboots
- End-to-end debugging across modules based on an existing connection
- Debugging ongoing connections

For more information about the troubleshooting connections, see Troubleshoot a connection, on page 439.

Troubleshoot a connection

This task enables you to identify and analyze connection issues by capturing and examining packet debug information for specific traffic flows.

Procedure

Step 1 Configure a filter to identify a connection using the **debug packet-condition** command.

Example:

```
Debug packet-condition match tcp 192.168.100.177 255.255.255.255 192.168.102.177
255.255.255.255
```

Step 2 Enable debugs for the interested modules and the corresponding levels. Enter the **debug packet** command.

Example:

```
Debug packet acl 5
```

Step 3 Start debugging the packets using the following command:

```
debug packet-start
```

Step 4 Fetch the debug messages from database to analyze the debug messages using the following command:

```
show packet-debugs
```

Step 5 Stop debugging the packets using the following command:

```
debug packet-stop
```

Advanced troubleshooting for the Secure Firewall Threat Defense device

You can use Packet Tracer and Packet Capture features to perform an in-depth troubleshooting analysis on a Secure Firewall Threat Defense device. A packet tracer allows a firewall administrator to inject a virtual packet into a security appliance and track the flow from ingress to egress. Along the way, the packet is evaluated against flow and route lookups, ACLs, protocol inspection, NAT, and intrusion detection. This utility is effective because it can simulate real-world traffic by specifying source and destination addresses with protocol and port information. Packet capture is available with the trace option, which provides you with a verdict as to whether the packet is dropped or successful.

For more information about the troubleshooting files, see Download advanced troubleshooting files, on page 434.

Packet capture

The packet capture feature with trace option allows real packets that are captured on the ingress interface to be traced through the system. The trace information is displayed at a later stage. These packets are not dropped on the egress interface, as they are real data-path traffic. Packet capture for Firewall Threat Defense devices supports troubleshooting and analysis of data packets.

Snort verdict processing

Once the packet is acquired, Snort detects the tracing flag that is enabled in the packet. Snort writes tracer elements, through which the packet traverses. Snort verdict as a result of capturing packets can be one of the following:

Table 28: Snort verdicts

Verdict	Description
Pass	Allow analyzed packet.
Block	Packet not forwarded.

Verdict	Description
Replace	Packet modified.
AllowFlow	Flow passed without inspection.
BlockFlow	Flow was blocked.
Ignore	Flow was blocked; occurs only for sessions with flows blocked on passive interfaces.
Retry	Flow is stalled, waiting on a enamelware or URL category/reputation query. In the event of a timeout, processing continues with an unknown result: in the case of enamelware, the file is allowed; in the case of URL category/reputation, AC rule lookup continues with an uncategorized and unknown reputation.

Based on the Snort verdict, the packets are dropped or allowed. For example, the packet is dropped if the Snort verdict is **BlockFlow**, and the subsequent packets in the session are dropped before reaching Snort. When the Snort verdict is **Block** or **BlockFlow**, the **Drop Reason** can be one of these:

Table 29: Drop reasons

Blocked or Flow Blocked by...	Cause
Snort	Snort is unable to process the packet, erg., snort can't decode packet since it is corrupted or has invalid format.
the App Id preprocessed	App Id module/preprocessed does not block packet by itself; but this may indicate that App Id detection causes other module (erg., firewall) to match a blocking rule.
the SSL preprocessed	There is a block/reset rule in SSL policy to match the traffic.
the firewall	There is a block/reset rule in firewall policy to match the traffic.
the captive portal preprocessed	There is a block/reset rule using the identity policy to match the traffic.
the safe search preprocessed	There is a block/reset rule using the safe-search feature in firewall policy to match the traffic.
the SI preprocessed	There is a block/reset rule a in Security Intelligence tab of AC Policy to block the traffic, erg., DNS or URL SI rule.
the filterer preprocessed	There is a block/reset rule in filterer tab of AC policy to match the traffic.

Blocked or Flow Blocked by...	Cause
the stream preprocessed	There is an intrusion rule blocking/reset stream connection, erg., blocking when TCP normalization error.
the session preprocessed	This session was already blocked earlier by some other module, so session preprocessed is blocking further packets of the same session.
the fragmentation preprocessed	Blocking because earlier fragment of the data is blocked.
the snort response preprocessed	There is a react snort rule, erg., sending a response page on a particular HTTP traffic.
the snort response preprocessed	There is a snort rule to send custom response on packets matching conditions.
the reputation preprocessed	Packet matches a reputation rule, erg., blocking a given IP address.
the x-Link2State preprocessed	Blocking due to buffer overflow vulnerability detected in SMTP.
back orifice preprocessed	Blocking due to detection of back orifice data.
the SMB preprocessed	There is a snort rule to block SMB traffic.
the file process preprocessed	There is file policy that blocks a file, erg., enamelware blocking.
the IPS preprocessed	There is a snort rule using IPS, erg., rate filtering.

The packet capture feature allows you to capture and download packets that are stored in the system memory. However, the buffer size is limited to 32 MB due to memory constraint. Systems capable of handling very high volume of packet captures exceed the maximum buffer size quickly and thereby the necessity of increasing the packet capture limit is required. It is achieved by using the secondary memory (by creating a file to write the capture data). The maximum supported file size is 10 GB.

When the **file-size** is configured, the captured data gets stored to the file and the file name is assigned based on the capture name **recapture**.

The **file-size** option is used when you need to capture packets with the size limit more than 32 MB.

For information, see the *Cisco Secure Firewall Threat Defense Command Reference*.

Use a packet capture trace

Packet capture is a utility that provides a live snapshot of network traffic passing the specified interface of a device based on a defined criteria. This process continues to capture the packets as long as it has not paused, or the allocated memory has not exhausted.

Packet capture data includes information from Snort and preprocessors about verdicts and actions the system takes while processing a packet. Multiple packet captures are possible at a time. You can configure the system to modify, delete, clear, and save captures.



Note Capturing packet data requires packet copy. This operation may cause delays while processing packets and may also degrade the packet throughput. We recommend that you use packet filters to capture specific traffic data.

Before you begin

To use the packet capture tool on Secure Firewall Threat Defense devices, you must be an Admin or Maintenance user.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Troubleshoot** > **Packet Capture**.

Step 2 Select a device.

Step 3 Click **Add Capture**.

Step 4 Enter the **Name** for capturing the trace.

Step 5 Select the **Interface** for the capturing the trace.

Step 6 Specify **Match Criteria** details:

a) Select the **Protocol**.
b) Enter the IP address for the **Source Host**.
c) Enter the IP address for the **Destination Host**.
d) Check **SGT number** check box, and enter a Security Group Tag (SGT).

Step 7 Specify **Buffer** details:

a) Enter a maximum **Packet Size**.
b) Enter a minimum **Buffer Size**.
c) Select either **Continuous Capture** if you want the traffic captured without interruption, or **Stop when full** if you want the capture to stop when the maximum buffer size is reached.

Note
If **Continues Capture** is enabled, and when the allocated memory is full, the oldest captured packets in the memory is overwritten by the new captured packets.

d) Check the check box of **Trace**, if you want to capture the details for each packet.
e) Enter the value in **Trace Count** field. Default value is 128 . You can enter values in the range of 1-1000.

Step 8 Click **Save**.

The packet capture screen displays the packet capture details and its status. To have the packet capture page auto refreshed, check the **Enable Auto Refresh** check box and enter the auto refresh interval in seconds.

You can do the following on the packet capture:

- **Edit** ( ) to modify the capture criteria.
- **Delete** ( ) to delete the packet capture and the captured packets.
- **Clear** ( ) to erase all the captured packets from a Packet Capture. To erase the captured packets from all of the existing packet captures, click **Clear All Packets**.
- **Pause** ( ) to temporarily halt capturing packets.
- **Save** ( ) to save a copy of captured packets on a local machine in ASCII or PCAP format. Choose the required format option, and click **Save**. The saved packet capture is downloaded to your local machine.
- To view the details of the packets being captured, click the required capture row.

Packet tracer tool overview

A packet tracer tool is a network analysis utility. It simulates packet flows by modeling source and destination addresses and protocol characteristics. The tool validates network policies by performing lookups against access rules, NAT, routing, and rate-limiting policies. It also captures and logs detailed tracing information for each packet to assist with policy troubleshooting and debugging.

Packet tracer tool simulation and PCAP replay

Use the packet tracer tool to simulate the complete data path of a packet, including slow-path and fast-path modules. This simulation helps verify whether your configured policies allow or deny intended traffic. This tool helps you confirm your policy effectiveness and troubleshoot when your packets are blocked.

You can use packet capture (PCAP) files to replay complete traffic flows. Make sure your PCAP file contains a single Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) flow with up to 100 packets. On replay, the tool initializes client and server entities and synchronizes packet transmission, logging trace results for each packet in the flow.

Run a packet trace

Follow these steps to run a packet trace on your Firewall Threat Defense device.

Before you begin

- Ensure you have Admin or Maintenance user privileges.
- Have any required PCAP (.pcap or .pcapng) files ready if you intend to upload or replay captured traffic.
- Confirm the device supports packet tracing.

Procedure

Step 1 In Cloud-Delivered Firewall Management Center, choose **Devices** > **Troubleshoot** > **Packet Tracer**.

Step 2 From the **Select Device** drop-down, choose the device where you want to run the trace.

Step 3 Choose one of the following options:

- **Use Protocol**: Configure protocol and parameters manually.

- **Upload or Edit a PCAP file**: Upload a packet capture (PCAP) file.

Step 4 To upload a PCAP file:

a) Click the **Upload or Edit a PCAP File** drop-down, and choose the **Upload a PCAP File** option. To use a recently uploaded file, click the file name from the available list.

Note
- Only .pcap and .pcapng file formats are supported.
- The file can contain up to 100 packets. All packets must be from the same Ethernet connection or from the same VLAN-encapsulated TCP or UDP connection.
- Multiflow PCAP files are not supported. Upload only single-flow PCAP files.

b) If you choose to upload a PCAP file, drag and drop the PCAP file into the **Upload PCAP** dialog box or browse and select the PCAP file. After you select the file, the upload process starts automatically.

Note
- After you upload the file, the **Protocol**, **Source Type**, and **Destination Type** fields will be grayed, and cannot be edited. To make changes to these fields, you must upload a new PCAP file.
- You can edit the source and destination IP addresses, source and destination ports, VLAN ID, destination MAC address (for a firewall in transparent mode), and the PCAP file name.

c) Go to Step 6.

Step 5 If you choose to perform a manual configuration, do the following:

a) From the **Ingress Interface** drop-down, choose the ingress interface for the packet trace.

Note
Do not select **VTI**; VTI as ingress is not supported.

b) From the **Protocol** drop-down, select the packet type and specify protocol characteristics:
- **ICMP**: Enter the ICMP type, code (0-255), and optionally, an identifier.
- **TCP**/**UDP**/**SCTP**: Enter the source and destination port numbers.
- **GRE**/**IPIP**: Enter the protocol number 0-255.
- **ESP**: Enter the Security Parameter Index (SPI) value for the source (0-4294967295).
- **RAWIP**: Enter the port number (0-255).

c) Select the **Source Type** and enter the source IP address.

Types include IPv4, IPv6, and FQDN (domain name, if using Cisco TrustSec)

d) Select the **Source Port**, **Destination Type**, and enter the destination IP address.

Destination type options vary depending on the source type that you select.

e) Select the **Destination Port**.

f) To trace packets to a subinterface, enter a **VLAN ID**.

This value is optional for non-subinterfaces only, since all the interface types can be configured in a subinterface.

g) Specify a **Destination MAC Address**.

If the Secure Firewall Threat Defense device is running in transparent firewall mode and the ingress interface is VTEP, you must provide a **Destination MAC Address** when you enter a **VLAN ID** value. However, if the interface is a bridge group member, you can provide a **Destination MAC Address** optionally when you enter a **VLAN ID** value, but you must provide it when you do not enter a **VLAN ID** value.

In routed firewall mode, **VLAN ID** and **Destination MAC Address** are optional if the input interface is a bridge group member.

h) To ignore security checks for the simulated packet, click **Bypass all security checks for simulated packet**. This enables packet tracer to trace the packet through the system, which would otherwise be dropped.

i) To send the packet out through the device's egress interface, click **Allow simulated packet to transmit from device**.

j) If you want the packet tracer to consider the injected packet as an IPsec/SSL VPN-decrypted packet, click **Treat simulated packet as IPsec/SSL VPN decrypt**.

Step 6 From the **Ingress Interface** drop-down, choose the ingress interface for the packet trace.

Note
Do not select VTI. VTI as ingress interface is not supported for packet tracer.

Step 7 To use a PCAP replay in the packet tracer, do the following:

a) Click **Select a PCAP File**.

b) To upload a new PCAP file, click **Upload a PCAP file**. To reuse a recently uploaded file, click the file name from the list.

Note
Only .pcap and .pcapng file formats are supported. The PCAP file can contain only a single TCP/UDP-based flow with a maximum of 100 packets. The maximum character limit for the PCAP file name (including the file formats) is 64.

c) In the **Upload PCAP** box, you can either drag a PCAP file or click to browse to the location where the file is stored, and select the file. When you select the file, the upload process starts automatically.

d) Go to this step

Step 8 To define the trace parameters, from the **Protocol** drop-down list, select the packet type for the trace, and specify the protocol characteristics:

- **ICMP** : Enter the ICMP type, ICMP code (0-255), and optionally, the ICMP identifier.
- **TCP** / **UDP** / **SCTP** : Enter the source and destination port numbers.
- **GRE** / **IPIP** : Enter the protocol number, 0-255.
- **ESP** : Enter the Security Parameter Index (SPI) value for the source. Valid range is 0-4294967295.
- **RAWIP** : Enter the port number. Valid range is 0-255.

Step 9 Select the **Source Type** for the packet trace, and enter the source IP address.

Source and destination types include IPv4, IPv6, and fully-qualified domain names (FQDN). You can specify IPv4 or IPv6 addresses and FQDN, if you use Cisco TrustSec.

Step 10 Select the **Source Port** for the packet trace.

Step 11 Select the **Destination** type for the packet trace, and enter the destination IP address.

Destination type options vary depending on the source type that you select.

Step 12 Select the **Destination Port** for the packet trace.

Step 13 Optionally, if you want to trace a packet where the Security Group Tag (SGT) value is embedded in the Layer 2 CMD header (TrustSec), enter a valid **SGT number** .

Step 14 If you want packet tracer to enter a parent interface, which is later redirected to a sub-interface, enter a **VLAN ID** .

This value is optional for non-sub-interfaces only, since all the interface types can be configured on a sub-interface.

Step 15 Specify a **Destination MAC Address** for the packet trace.

If the Secure Firewall Threat Defense device is running in transparent firewall mode, and the ingress interface is VTEP, **Destination MAC Address** is required if you enter a value in **VLAN ID**. Whereas if the interface is a bridge group member, **Destination MAC Address** is optional if you enter a **VLAN ID** value, but required if you do not enter a **VLAN ID** value.

If the Secure Firewall Threat Defense is running in routed firewall mode, **VLAN ID** and **Destination MAC Address** are optional if the input interface is a bridge group member.

Step 16 (Optional) If you want the packet-tracer to ignore the security checks on the simulated packet, click **Bypass all security checks for simulated packet** . This enables packet-tracer to continue with tracing of packet through the system which, otherwise would have been dropped.

Step 17 (Optional) To allow the packet to be sent out through the egress interface from the device, click **Allow simulated packet to transmit from device** .

Step 18 (Optional) If you want the packet-tracer to consider the injected packet as an IPsec/SSL VPN decrypted packet, click **Treat simulated packet as IPsec/SSL VPN decrypt** .

Step 19 Click **Trace**.

Step 20 (Optional) If you modify any values, ensure you click **Save PCAP** to save your changes, then begin the trace.

Step 21 (Optional) If you do not save the modified values of the PCAP file and click **Trace**, the **Unsaved PCAP changes** dialog box appears and prompts you to save the file.

a. Check the **Save PCAP file** check box.

b. Enter a name for the PCAP file in the **Name** field.

c. Click **Save and Trace** to save the changes and continue with packet trace.

Note
The PCAP file name changes to the name entered in the **Name** field.

Step 22 You can track the status of trace in the **Events & Logs** > **Analysis** > **Audit Logs** window. These tasks can be tracked:

- Saving a PCAP file

- Uploading a PCAP file
- Details of the packet trace

The **Trace Result** displays the results for each phase that the PCAP packets have traveled through the system. Click an individual packet to view the traces results for the packet. You can do the following:

- Copy (Copy) the trace results to the clipboard.
- Expand or collapse (Expand or collapse) the displayed results.
- Maximize (Maximize) the trace result window.

The time elapsed information, useful to gauge the processing efforts, is displayed for each phase. The results section also displays the total time taken for packets flowing from an ingress to an egress interface.

The **Trace History** pane displays the stored trace details for each PCAP trace. It can store up to 100 packet traces. You can select a saved trace and run the packet trace activity again. You can do the following:

- Search for a trace using any of the trace parameters.
- Disable saving of the trace to history using the Slider button.
- Delete specific trace results.
- Clear all the traces.

CPU profiler overview

The CPU profiler gathers data on CPU usage by individual modules or inspectors of Snort 3, while packets are being processed within a specific timeframe. It provides information on the amount of CPU time that each module consumes in relation to the total CPU usage of the Snort 3 process. Using the CPU profiler does not require reloading the configuration or restarting Snort 3, thus minimizing downtime. The profiling result displays the processing time taken by all the modules during the last profiling session. CPU profiling results are saved in JSON format in the threat defense device and synchronized in the management center.

Caution CPU profiling might cause system performance degrade of approximately 3 percent.

Use the CPU profiler

Use the CPU profiler to identify which Snort 3 modules or inspectors are consuming the most processing time and contributing to performance issues.

Before you begin

You must have Version 7.6 or later devices with Snort 3 to use CPU Profiling.

Procedure

Step 1 From the management center, choose **Devices** > **Troubleshoot** > **Snort 3 Profiling**.

Step 2 Click the **CPU Profiling** tab.

Step 3 From the **Select device for CPU Profiling** drop-down list, choose a device for CPU profiling.

Note
You can run multiple profiling sessions simultaneously on different devices.

Step 4 To start a CPU profiling session, click **Start**. (The session automatically stops after 120 minutes.)

You can stop the profiling session at any time by clicking **Stop**. However, canceling it before the scheduled 120 minutes might not give precise results.

Note
While the CPU profiling session is in progress, a task is created. Click **Notifications** (message center), and then click the **Tasks** tab to view the details.

The latest profiling result is automatically displayed in the **CPU Profiling Results** section. The table contains statistics for the processing time taken by all the Snort 3 modules or inspectors during the last profiling session. You can view the CPU profiler output in tabular format:

- **Module**: Name of the module or inspector.
- **% Total of CPU Time**: Percentage of time taken by the module relative to the overall time taken by Snort 3 for processing traffic. If this value is considerably higher than that of other modules, it means the module is contributing more to the unsatisfactory performance of Snort 3.
- **Time (µs)**: Total time (in microseconds) taken by each module.
- **Avg/Check**: Average time taken by the module each time the module was invoked.
- **% Caller**: Time taken by the submodule (if configured) relative to the main module. This value is used for debugging purposes.

Step 5 (Optional) Click **Download Snapshot** to download the profiling result. The downloaded file is in CSV format and contains all the fields from the profiling results page.

Step 6 (Optional) Click the **Filter by % of Snort time** toggle button to filter out modules whose execution consumed more than *n*% of the profiling time.

Step 7 (Optional) Use the **Search** field to search through the fields in the **CPU Profiling Results** table.

Note
Except for **Module**, you can click any of the other column headers to sort the data.

Step 8 (Optional) Click the **Profiling History** section (collapsible panel on the left) to expand it and view a set of cards representing the previous profiling sessions for a chosen device. When you click a card from history, the details are displayed in the **CPU Profiling Results** section.

The CPU profiling session provides detailed statistics about the processing time of each Snort 3 module, helping you identify performance bottlenecks.

Note If you initiate a deployment while CPU profiling is running, the profiling session is automatically terminated to accommodate the deployment, except for deployments resulting from changes to the access control policy rules and security intelligence. You must run the CPU profiling for the device again.

Rule profiler overview

The Snort 3 rule profiler gathers data on the time spent processing a set of Snort 3 intrusion rules, highlights potential issues, and shows the rules with unsatisfactory performance. Within the Snort 3 engine, the rule profiler inspects traffic using the Snort 3 intrusion detection mechanism. The rule profiler displays the top 100 IPS rules that required the maximum time to check. Triggering the rule profiler does not require Snort 3 reload or restart. The rule profiling results are saved in JSON format in the threat defense device and synchronized in the management center.

Use the rule profiler

The Rule Profiler helps you identify IPS rules that consume excessive processing time, allowing you to optimize performance by analyzing rule execution statistics and identifying poorly performing rules that may need adjustment or replacement.

Before you begin

You must have version 7.6 or later devices with Snort 3 to use Rule Profiling.

Procedure

Step 1 From the management center, choose **Devices** > **Troubleshoot** > **Snort 3 Profiling**.

Step 2 Click the **Rule Profiling** tab.

Step 3 From the **Select device for Rule Profiling** drop-down list, choose a device for rule profiling.

Note
You can run multiple profiling sessions simultaneously on different devices.

Step 4 To start a rule profiling session, click **Start**.

The session automatically stops after 120 minutes. You can stop the profiling session at any time by clicking **Stop**. However, canceling it before the scheduled 120 minutes might not give precise results.

Note
While the rule profiling session is in progress, a task is created. Click **Notifications** > **Tasks** to view the details.

The latest profiling result is automatically displayed in the **Rule Profiling Results** section. The table contains statistics for the rules that took the most time to process (sorted in descending order) by the total time (in microseconds [µs]). You can view the IPS rule profiler output in tabular format:

- **% of Snort Time**: Time spent for processing the rule, relative to the total time of the Snort 3 operation.
- **Rev**: Revision number of the rule.

- **Checks**: Number of times the IPS rule was executed.
- **Matches**: Number of times the IPS rule fully matched in the traffic.
- **Alerts**: Number of times the IPS rule triggered an IPS alert.
- **Time**: Time (in microseconds) spent by Snort in checking the IPS rule.
- **Avg/Check**: Average time spent by Snort on a single check of the rule.
- **Avg/Match**: Average time spent by Snort on a single check, which resulted in a match.
- **Avg/Non-Match**: Average time spent by Snort on a single check, which did not result in a match.
- **Timeouts**: Number of times the rule exceeded the Rule Handling - Threshold configured in the Latency-Based Performance Settings of the access control policy.
- **Suspends**: Number of times the rule was suspended due to consecutive threshold violations.

Step 5 Click **Download Snapshot** to download the profiling result.

The downloaded file is in CSV format and contains all the fields from the profiling results page.

Step 6 Click the **Filter by % of Snort time** toggle button to filter out rules whose execution took more than *n*% of the profiling time.

In general, a rule is considered as performing unsatisfactorily if it consumes 0.2% or more of Snort's overall processing time.

Step 7 Use the **Search** field to search through the fields in the **Rule Profiling Results** table.

Step 8 Click the **Rule Profiling History** section (collapsible panel on the left) to expand it and view a set of cards representing the previous profiling sessions for a chosen device.

When you click a card from the history, the details are displayed in the **Rule Profiling Results** section.

Note If you initiate a deployment while rule profiling is running, the profiling session is automatically terminated to accommodate the deployment, except for deployments resulting from changes to the access control policy rules and security intelligence. You must run the rule profiling for the device again.

Use the Firewall Threat Defense diagnostic CLI from the web interface

Use the diagnostic CLI to execute selected commands for quick troubleshooting of your device from the web interface without requiring direct SSH access to the device.

You can execute the selected Firewall Threat Defense diagnostic CLI commands from the Cloud-Delivered Firewall Management Center. The commands **ping** (except **ping system**), **traceroute**, and select **show** commands run in the diagnostic CLI rather than the regular CLI.

When you run the **show** commands, if the message `Unable to execute the command properly. Please see logs for more details` is displayed, it means that the command is not valid in the diagnostic CLI. For example, **show access-list** works, but this message will be displayed if you enter **show**

access-control-policy. To use non-diagnostic commands, use SSH to log in to a device outside management center.

For more information on the Firewall Threat Defense CLI, see the *Cisco Secure Firewall Threat Defense Command Reference*.

Before you begin

- You must be an Admin, Maintenance, or Security Analyst to use the diagnostic CLI.
- The purpose of diagnostic CLI is to enable the quick use of a few commands that are useful in troubleshooting a device. For access to the full range of commands, open an SSH session directly with the device.
- In deployments using Cloud-Delivered Firewall Management Center high availability,diagnostic CLI is available only in the active Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Troubleshoot** > **Threat Defense CLI**.

You can also access the CLI tool through the health monitor for the device (**System** (⚙) > **Health** > **Monitor**). From there, you can select the device, click the **View System and Troubleshoot Details** link, click **Advanced Troubleshooting**, then click **Threat Defense CLI** on that page.

Step 2 From the **Device** drop-down list, choose the device on which to execute the diagnostic command.

Step 3 From the **Command** drop-down list, choose the command that you want to execute.

Step 4 Enter the command parameters in the **Parameters** field.

See the *Cisco Secure Firewall Threat Defense Command Reference* for the valid parameters.

For example, to execute **show access-list** command, choose **show** from the **Command** drop-down list, then enter **access-list** in the **Parameters** field.

Note
Do not type the full command in the **Parameters** field. Type only the relevant keywords.

Step 5 Click **Execute** to view the command output.

If the message `Unable to execute the command properly. Please see logs for more details.` is displayed, examine the parameters closely. There might be syntax errors.

This message can also mean that the command you are trying to execute is not a valid command within the context of the diagnostic CLI (which you have accessed from the device using the **system support diagnostic-CLI** command). Log in to the device using SSH to use these commands.

Feature-specific troubleshooting

This reference provides a comprehensive list of troubleshooting resources organized by feature. Each entry includes links to detailed troubleshooting information, best practices, and configuration guidance for specific system features and components.

Table 30: Feature-specific troubleshooting topics

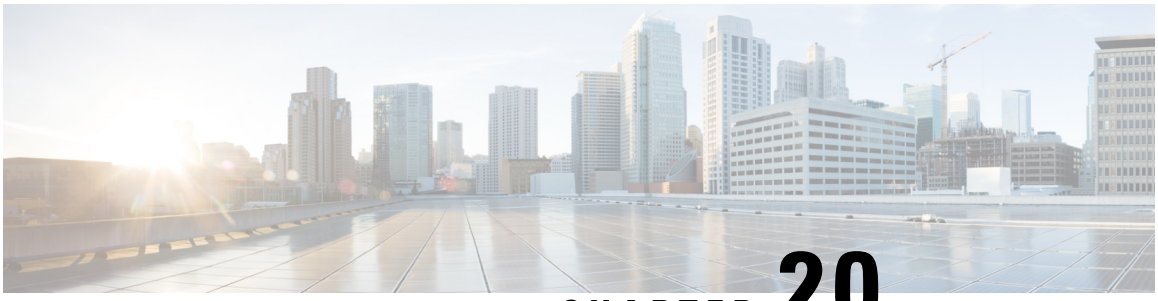
Feature	Relevant Troubleshooting Information
Application control	*Best Practices for Application Control* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
LDAP external authentication	Troubleshoot LDAP authentication connections, on page 261
Licensing	Troubleshoot Smart Licensing, on page 299
User rule conditions	*Troubleshoot User Control* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
User identity sources	For troubleshooting information on ISE/ISE-PIC, TS Agent Identity Source, Captive Portal Identity Source, and Remote Access VPN Identity Source, see the corresponding sections in the *Cisco Secure Firewall Management Center Device Configuration Guide* Troubleshoot LDAP authentication connections, on page 261
URL filtering	*Troubleshoot URL Filtering* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Realms and user data downloads	*Troubleshoot Realms and User Downloads* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Custom Security Group Tag (SGT) rule conditions	*Custom SGT Rule Conditions* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
SSL rules	Chapter on SSL rules in the *Cisco Secure Firewall Device Manager Configuration Guide*
Cisco Threat Intelligence Director (TID)	*Troubleshoot Secure Firewall Threat Intelligence Director* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Secure Firewall Threat Defense syslog	*About Configuring Syslog* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Intrusion performance statistics	*Intrusion Performance Statistic Logging Configuration* in the *Cisco Secure Firewall Management Center Device Configuration Guide*
Connection-based Troubleshooting	Connection-based troubleshooting, on page 439



PART VII

Tools

- Backup/Restore, on page 457
- Scheduling, on page 471
- Import/Export, on page 477



CHAPTER 20

Backup/Restore

- About Backup and Restore, on page 457
- Requirements for Backup and Restore, on page 458
- Guidelines and Limitations for Backup and Restore, on page 459
- Best Practices for Backup and Restore, on page 460
- Back Up Managed Devices, on page 462
- Restore Security Cloud Control-Managed Devices, on page 464

About Backup and Restore

The ability to recover from a disaster is an essential part of any system maintenance plan. As part of your disaster recovery plan, we recommend that you perform periodic backups.

On-Demand Backups

You can perform on-demand backups for Secure Firewall Threat Defense devices in Security Cloud Control.

For more information, see Back Up Managed Devices, on page 462.

Scheduled Backups

You can use the scheduler on Cloud-Delivered Firewall Management Center to automate backups. You can also schedule remote device backups from Security Cloud Control.

The Cloud-Delivered Firewall Management Center setup process schedules weekly configuration-only backups, to be stored in Cisco's secure cloud storage. This is not a substitute for full off-site backups. After initial setup finishes, you must review your scheduled tasks and adjust them to fit your organization's needs.

For more information, see Scheduling, on page 471.

Store Backup Files

When you back up a device, the Cloud-Delivered Firewall Management Center stores the backup files in its secure cloud storage.

Note Cloud-Delivered Firewall Management Center stores only the 5 most recent Firewall Threat Defense backups per device.

To know how to download an existing backup file, see Step 3 in Restore a Firewall Threat Defense Device, on page 464.

For more information, see Back Up Managed Devices, on page 462.

Restore Managed Devices

You must use the Firewall Threat Defense CLI to restore the Firewall Threat Defense device.

For more information, see Restore Security Cloud Control-Managed Devices, on page 464.

What Is Backed Up?

Device backups are always configuration-only.

What Is Restored?

Restoring configurations overwrites *all* backed-up configurations.

Make sure you understand and plan for the following:

- You cannot restore what is not backed up.
- The Firewall Threat Defense restore process removes VPN certificates and all VPN configurations from Firewall Threat Defense devices, including certificates added after the backup was taken. After you restore a Firewall Threat Defense device, you must re-add/re-enroll all VPN certificates, and redeploy the device.

Requirements for Backup and Restore

Backup and restore have the following requirements:

Model Requirements: Backup

You can back up:

- Firewall Threat Defense standalone devices, native instances, container instances, HA pairs, and clusters
- Firewall Threat Defense Virtual for VMware devices, either standalone or HA pairs, and clusters

Backup is *not* supported for:

- Firewall Threat Defense Virtual implementations *other than* for VMware

If you must replace a device where backup and restore is not supported, you must manually recreate device-specific configurations.

Model Requirements: Restore

A replacement managed device must be the same model as the one you are replacing, with the same number of network modules and same type and number of physical interfaces.

Version Requirements

As the first step in any backup, note the patch level. To restore a backup, the old and the new appliance must be running the same Firewall version, including patches.

License Requirements

Address licensing or orphan entitlement concerns as described in the best practices and procedures. If you notice licensing conflicts, contact Cisco TAC.

Guidelines and Limitations for Backup and Restore

Backup and restore have the following guidelines and limitations.



Caution Users with CLI access can access the Linux shell access with the **expert** command, which can present a security risk. For system security reasons, we strongly recommend:

- Only use the Linux shell under TAC supervision or when explicitly instructed by Firewall and Security Cloud Control user documentation.
- Restrict the list of users with Linux shell access..
- Do not add users directly in the Linux shell; only use the procedures in this chapter.

Backup and Restore Is for Disaster Recovery/Return Material Authorization

Backup and restore are primarily intended for Return Material Authorization (RMA) scenarios. Before you begin the restore process of a faulty or failed physical appliance, contact for replacement hardware.

Backup and Restore Is Not Configuration Import/Export

A backup file contains information that uniquely identifies an appliance, and cannot be shared. Do not use the backup and restore process to copy configurations between appliances or devices, or as a way to save configurations while testing new ones. Instead, use the import/export feature.

For example, Firewall Threat Defense device backups include the device's management IP address and all information the device needs to connect to its managing Security Cloud Control. Do not restore an FTD backup to a device being managed by a different manager; the restored device attempts to connect to the manager specified in the backup.

Restore Is Individual and Local

You restore threat defense devices individually and locally. This means:

- You cannot batch-restore to high availability (HA) devices. The restore procedures in this guide explain how to restore in an HA environment.

- You cannot use Security Cloud Control to restore a device. For Firewall Threat Defense devices, you must use the Firewall Threat Defense CLI, except for the ISA 3000 zero-touch restore, which uses an SD card and the reset button.
- You cannot use a Security Cloud Control user account to log into and restore one of its managed devices. Security Cloud Control and Firewall Threat Defense devices maintain their own user accounts.

Best Practices for Backup and Restore

Backup and restore have the following best practices.

When to Back Up

We recommend backing up during a maintenance window or other time of low use.

You must back up in the following situations:

- Regular backups.

  As part of your disaster recovery plan, we recommend that you perform periodic backups.

- Before upgrade or reimage.

  If an upgrade fails catastrophically, you may have to reimage and restore. Reimaging returns most settings to factory defaults, including the system password. If you have a recent backup, you can return to normal operations more quickly.

- After upgrade.

  Ensure that you back up the device after upgrade, so that you have a freshly upgraded device backup.

Maintaining Backup File Security

Backup files are stored as unencrypted archive (.tar) files; they should be stored in a secure repository.

Private keys in PKI objects, which represent the public key certificates and paired private keys that are required to support your deployment are decrypted before they are backed up. The keys are reencrypted with a randomly generated key when you restore the backup.

Backup file should be stored securely.

Backup and Restore in Firewall Threat Defense High Availability Deployments

In a Firewall Threat Defense HA deployment, you must:

- Back up the device pair from the Cloud-delivered Firewall Management Center, but restore individually and locally from the Firewall Threat Defense CLI.

  The backup process produces unique backup files for Firewall Threat Defense HA devices. Do not restore one HA peer with the backup file from the other. A backup file contains information that uniquely identifies an appliance, and cannot be shared.

  A Firewall Threat Defense HA device's role is noted in its backup file name. When you restore, make sure you choose the appropriate backup file: primary vs secondary.

- Do *not* suspend or break HA before you restore.

Maintaining the HA configuration ensures that replacement devices can easily reconnect after restore. Note that you will have to resume HA synchronization to make this happen.

- Do *not* run the restore CLI command on both peers at the same time.

  Assuming you have successful backups, you can replace either or both peers in an HA pair. Any physical replacement tasks that you can perform simultaneously: unracking, reracking, and so on. However, do *not* run the restore command on the second device until the restore process completes for the first device, including the reboot.

Backup and Restore in Firewall Threat Defense Clustering Deployments

In the Firewall Threat Defense clustering deployment, you must:

- Back up the entire cluster from the Cloud-delivered Firewall Management Center, but restore nodes individually and locally from the threat defense CLI.

  The backup process produces a bundled tar file that includes unique backup files for each cluster node. Do not restore one node with the backup file from another. A backup file contains information that uniquely identifies a device, and cannot be shared.

  The node's role is noted in its backup file name. When you restore, make sure you choose the appropriate backup file: control or data.

  You cannot back up individual nodes. If a data node fails to back up, the management center will still back up all other nodes. If the control node fails to back up, the backup is canceled.

- All the nodes that are part of the cluster must be registered in the management center, for the backup to be successful.

- Do *not* suspend or break clustering before you restore. Maintaining the cluster configuration ensures replacement devices can easily reconnect after restore.

- Do *not* run the **restore** CLI command on multiple nodes at the same time. We recommend that you restore the control node first and wait until it rejoins the cluster before you restore any data nodes.

  Assuming you have successful backups, you can replace multiple nodes in the cluster. Any physical replacement tasks you can perform simultaneously: unracking, reracking, and so on. However, do *not* run the **restore** command on an additional node until the restore process completes for the previous node, including the reboot.

Before Backup

Ensure that any firewall located between the Firewall Threat Defense and Security Cloud Control is configured to allow outbound HTTPS traffic from the Firewall Threat Defense to the Security Cloud Control host.

Before Restore

Before restore, you must:

- Revert licensing changes.

  Revert any licensing changes made since you took backup.

  Otherwise, you may have license conflicts or orphan entitlements after the restore. However, do *not* unregister from Cisco Smart Software Manager (CSSM). If you unregister from CSSM, you must unregister again after you restore, then re-register.

After the restore completes, reconfigure licensing. If you notice licensing conflicts or orphan entitlements, contact Cisco TAC.

- Disconnect faulty appliances.

  Disconnect the management interface, and for devices, the data interfaces.

  Restoring a Firewall Threat Defense device sets the management IP address of the replacement device to the management IP address of the old device. To avoid IP address conflicts, disconnect the old device from the management network before you restore the backup on its replacement.

- Do *not* unregister managed devices.

  Whether you are restoring a managed device, do not unregister devices from the Security Cloud Control, even if you physically disconnect an appliance from the network.

  If you unregister, you must redo some device configurations, such as security zone to interface mappings. After you restore, Security Cloud Control and devices should begin communicating normally.

- Reimage.

  In an RMA scenario, the replacement appliance arrives configured with factory defaults. However, if the replacement appliance is already configured, we recommend you reimage. Reimaging returns most settings to factory defaults, including the system password. You can only reimage to major versions, so you must patch after you reimage.

  If you do not reimage, keep in mind that Security Cloud Control intrusion events and file lists are merged rather than overwritten.

After Restore

After restore, you must:

- Reconfigure anything that was not restored.

  This can include reconfiguring licensing and audit log server certificate settings. You also must re-add/re-enroll failed Firewall Threat Defense VPN certificates.

- Deploy.

  After you restore a device, deploy to that device. You *must* deploy. If the device or devices are not marked out of date, force deploy from the Device Management page.

Back Up Managed Devices

You can perform on-demand or scheduled backups for supported Secure Firewall Threat Defense devices using Cloud-Delivered Firewall Management Center without a backup profile.

For more information, see Back Up a Firewall Threat Defense Device from Cloud-delivered Firewall Management Center, on page 463.

Related Information

Back Up a Firewall Threat Defense Device from Cloud-delivered Firewall Management Center

Use this procedure to perform an on-demand backup of any of the following devices:

- Firewall Threat Defense: Physical devices, standalone, HA, or cluster
- Firewall Threat Defense Virtual: VMware, standalone, HA, or cluster

Backup and restore is not supported for any other platforms or configurations.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. Do not skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 458
- Guidelines and Limitations for Backup and Restore, on page 459
- Best Practices for Backup and Restore, on page 460



Caution Users with CLI access can access the Linux shell access with the **expert** command, which can present a security risk. For system security reasons, we strongly recommend:

- Only use the Linux shell under TAC supervision or when explicitly instructed by Firewall and Security Cloud Control user documentation.
- Restrict the list of users with Linux shell access..
- Do not add users directly in the Linux shell; only use the procedures in this chapter.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 From the Security Cloud Control menu, navigate **Administration** > **Firewall Management Center** to open the **Services** page.

Step 3 Select **Cloud-Delivered FMC** and in the **Management** pane, click **Devices** to navigate to the Cloud-Delivered Firewall Management Center user interface.

Step 4 Select **Administration** > **Advanced** > **Backup & Restore**
The **Device backup** page displays the device name, backup file name, device model, device version, and the backup file size.

Step 5 Click **Create device backup**.

Step 6 Select one or more Firewall Threat Defense devices in **Select devices to backup**.

You can filter the list based on the device groups, or device type—Standalone, High Availability, and Cluster. For clustering, choose the cluster. You cannot perform backups on individual nodes.

Note
If the device name has special characters other than hyphen, period, and underscore, those special characters will be replaced with an underscore.

Step 7 Click **Next** to view the summary of your selections.

Step 8 Click **Start backup** to start the on-demand backup.

Step 9 Monitor the progress under **Tasks** in the **Notifications** pane.

- The notifications will display one consolidated entry per backup action.
- To view detailed information for each of the devices being backed up, click **Details**.
- In the **Device backup information** dialog, you can monitor the backup progress at granular level.
- For cluster backup, you can view node-wise backup progress.
- To download the backup report in PDF format, click **Download report**.

Restore Security Cloud Control-Managed Devices

For Firewall Threat Defense devices, you must use the Firewall Threat Defense CLI to restore from backup. You cannot use the Cloud-Delivered Firewall Management Center to restore a device.

Related Information:

Restore a Firewall Threat Defense Device

Firewall Threat Defense backup and restore is intended for RMA. Restoring the configurations overwrites *all* configurations on the device, including the management IP address. It also reboots the device.

In case of hardware failure, this procedure outlines how to replace a firewall device, either standalone or in an HA pair. It assumes you have access to a successful backup of the device or devices you are replacing.



Note Do *not* unregister from the Security Cloud Control, even when disconnecting a device from the network. In a Firewall Threat Defense HA deployment, do *not* suspend or break HA. Maintaining these links ensures replacement devices can automatically reconnect after a restore.

In a Firewall Threat Defense HA deployment, you can use this procedure to replace either or both peers. To replace both, perform all steps on both devices simultaneously, except the restore CLI command itself. You cannot replace a Firewall Threat Defense HA device without a successful backup.

Note Starting from the next upgrade of Cloud-Delivered Firewall Management Center, Cisco will only store the 5 most recent Firewall Threat Defense backups per device and purge all the other backups. If you require the other backups for compliance purposes, we recommend that you download them to your private environment by May 31, 2025.

To know how to download an existing backup file, see Step 3 in the procedure below.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. Do not skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 458
- Guidelines and Limitations for Backup and Restore, on page 459
- Best Practices for Backup and Restore, on page 460

Procedure

Step 1 Contact Cisco TAC for replacement hardware.

Obtain an identical model, with the same number of network modules and same type and number of physical interfaces. You can begin the RMA process from the Cisco Returns Portal.

Step 2 Select **Administration** > **Advanced** > **Backup & Restore**
The **Device backup** page displays the device name, backup file name, device model, device version, and the backup file size.

Step 3 Locate a successful backup of the faulty device from **Device backup**.

Use **Download** that downloads the backup file(s) to your local storage or **Export backup link** that generates a URL to download the backup and exports it to a CSV file that gets downloaded. Use the URL to download the backup to a secure location. Note that the URL is valid only for six hours, after which you must export it again to get a different URL. In a cluster device backup, you can now download a specific cluster node backup.

In a Firewall Threat Defense HA deployment, you back up the pair as a unit but the backup process produces unique backup files for each device in the pair. The device's role is noted in the backup file name.

If the only copy of the backup is on the faulty device, copy it somewhere else now. If you reimage the device, the backup will be erased. If something else goes wrong, you may not be able to recover the backup.

The replacement device will need the backup, but can retrieve it with the secure copy (SCP) command during the restore process. We recommend you put the backup somewhere SCP-accessible to the replacement device. Or, you can copy the backup to the replacement device itself.

Step 4 Remove (unrack) the faulty device and disconnect all interfaces. In threat defense HA deployments, this includes the failover link.

See the hardware installation and getttign started guides for your model: Cisco Firepower NGFW: Install and Upgrade Guides.

Note

Do not unregister from the management center, even when disconnecting a device from the network. In threat defense HA deployments, do not suspend or break HA. Maintaining these links ensures replacement devices can automatically reconnect after restore.

Step 5 Install the replacement device and connect it to the management network.

Connect the device to power and the management interface to the management network. In Firewall Threat Defense HA deployments, connect the failover link. However, do *not* connect the data interfaces.

See the hardware installation guide for your model: Cisco Firepower NGFW: Install and Upgrade Guides.

Step 6 (Optional) Reimage the replacement device.

In an RMA scenario, the replacement device will arrive configured with factory defaults. If the replacement device is not running the same major version as the faulty device, we recommend you reimage.

See the Cisco Secure Firewall ASA and Threat Defense Reimage Guide.

Step 7 Perform initial configuration on the replacement device.

Access the Firewall Threat Defense CLI as the admin user. You can use the console or you can SSH to the factory-default management interface IP address (192.168.45.45). A setup wizard prompts you to configure the management IP address, gateway, and other basic network settings.

See the initial configuration topics in the getting started guide for your model: Cisco Firepower NGFW: Install and Upgrade Guides.

Note
If you need to patch the replacement device, start the management center registration process as described in the getting started guide. If you do *not* need to patch, do not register.

Step 8 Make sure the replacement device is running the same Firewall software version, including patches, as the faulty device.

The existing device should not be deleted from the management center. The replacement device should be unmanaged from the physical network and the new hardware as well as the replacing Firewall Threat Defense patch should have the same version. The Firewall Threat Defense CLI does not have an upgrade command. To patch:

a) From the management center web interface, complete the device registration process: See *Add a Device to the Management Center* in Cisco Secure Firewall Management Center Device Configuration Guide.

   Create a new AC policy and use the default action "Network Discovery". Leave this policy as is; do not add any features or modifications. This is being used to register the device and deploy a policy with no features so that you do not require licenses, and you will then be able to patch the device. Once backup is restored, it should restore the licensing and policy into the expected state.

b) Patch the device: Cisco Firewall Management Center Upgrade Guide.

c) Unregister the freshly patched device from the management center: See *Delete a Device from the Management Center* in Cisco Secure Firewall Management Center Device Configuration Guide.

   If you do not unregister, you will have a ghost device registered to the management center after the restore process brings your "old" device back up.

Step 9 Make sure the replacement device has access to the backup file.

The restore process can retrieve the backup with SCP, so we recommend you put the backup somewhere accessible. Or, you can manually copy the backup to the replacement device itself, to `/var/sf/backup`.

Step 10 From the FTD CLI, restore the backup.

Access the Firewall Threat Defense CLI as the admin user. You can use the console or you can SSH to the newly configured management interface (IP address or hostname). Keep in mind that the restore process will change this IP address.

To restore:

- With SCP: **restore remote-manager-backup location scp-hostname username filepath backup tar-file**
- From the local device: **restore remote-manager-backup backup tar-file**

Step 11 Log in to Security Cloud Control and wait for the devices to connect.

When the restore is done, the device logs you out of the CLI, reboots, and automatically connects to Security Cloud Control. At this time, the device should appear out of date.

At this time, the device should appear out of date.

Step 12 Before you deploy, perform any post-restore tasks and resolve any post-restore issues:

- Resolve licensing conflicts or orphan entitlements. Contact Cisco TAC.
- Resume HA synchronization.
- Re-add/re-enroll all VPN certificates. The restore process removes VPN certificates from FTD devices, including certificates added after the backup was taken.

Step 13 Deploy configurations.

You must deploy. If a restored device is not marked out of date, force deploy from the Device Management page.

Step 14 Connect the device's data interfaces.

See the hardware installation guide for your model: Cisco Secure Firewall Threat Defense: Install and Upgrade Guides.

Restore Firewall Threat Defense from Backup: Firewall Threat Defense Virtual

Use this procedure to replace a faulty or failed Firewall Threat Defense Virtual device for VMware.

In Firewall Threat Defense HA and clustering deployments, you can use this procedure to replace all peers. To replace all, perform all steps on all devices simultaneously, except the **restore** CLI command itself.



Note Do *not* unregister from the Firewall Management Center, even when disconnecting a device from the network. In Firewall Threat Defense HA and clustering deployments, do *not* suspend or break high availibility or clusters. Maintaining registration ensures that replacement devices can automatically reconnect after restore.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. Do not skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 458
- Guidelines and Limitations for Backup and Restore, on page 459
- Best Practices for Backup and Restore, on page 460

Procedure

Step 1 Navigate **System(⚙) > Tools > Backup/Restore**.

Step 2 Locate a successful backup of the faulty device from **Device Backups** under **Backup Management**.

For clustering, node backup files are bundled together in a single compressed file for the cluster (*cluster_name.timestamp***.tar.gz**). Before you can restore nodes, you need to extract the individual node backup files (*node_name***_control_***timestamp***.tar** or *node_name***_data_***timestamp***.tar**).

Use **Download** that downloads the backup file(s) to your local storage or **Export Backup Links** that generates a URL to download the backup and exports it to a CSV file that gets downloaded. Use the URL to download the backup to a secure location. Note that the URL is valid only for six hours, after which you must export it again to get a different URL.

In Firewall Threat Defense HA deployments, you back up the pair as a unit, but the backup process produces unique backup files for each device in the pair. The device's role is noted in the backup file name.

If the only copy of the backup is on the faulty device, copy it somewhere else now. If you reimage the device, the backup will be erased. If something else goes wrong, you may not be able to recover the backup.

The replacement device needs the backup, but can retrieve it with SCP during the restore process. We recommend you put the backup somewhere SCP-accessible to the replacement device. Or, you can copy the backup to the replacement device itself.

Step 3 Remove the faulty device.

Shut down, power off, and delete the virtual machine. For procedures, see the documentation for your virtual environment.

Step 4 Deploy a replacement device.

See the Cisco Firepower Threat Defense Virtual for VMware Getting Started Guide.

Step 5 Perform initial configuration on the replacement device.

Use the VMware console to access the Firewall Threat Defense Virtual CLI as the admin user. A setup wizard prompts you to configure the management IP address, gateway, and other basic network settings.

Do not set the same management IP address as the faulty device. This can cause problems if you need to register the device in order to patch it. The restore process will correctly reset the management IP address.

See the CLI setup topics in the getting started guide: Cisco Firepower Threat Defense Virtual for VMware Getting Started Guide.

Step 6 Make sure that the replacement device is running the same Firewall software version, including patches, as the faulty device.

Ensure that the existing device should not be deleted from the Security Cloud Control. The replacement device should be unmanaged from the physical network and the new hardware and the replacing Firewall Threat Defense Virtual patch should have the same version. The Firewall Threat Defense Virtual CLI does not have an upgrade command. To patch:

a. Complete the Firewall Threat Defense Virtual registration process in Security Cloud Control.

b. Patch the Firewall Threat Defense Virtual device.

c. Unregister the freshly patched device from Security Cloud Control.

Step 7 Make sure that the replacement device has access to the backup file.

The restore process can retrieve the backup with SCP, so we recommend you put the backup somewhere accessible. Or, you can manually copy the backup to the replacement device itself, to `/var/sf/backup`. For clusters, make sure you extract the individual node backup file from the main cluster bundle.

Step 8 From the Firewall Threat Defense CLI, restore the backup.

Access the Firewall Threat Defense Virtual CLI as the admin user. You can use the console or you can SSH to the newly configured management interface (IP address or hostname). Keep in mind that the restore process will change this IP address.

To restore:

- With SCP: **restore remote-manager-backup location scp-hostname username filepath backup tar-file**
- From the local device: **restore remote-manager-backup backup tar-file**

In Firewall Threat Defense HA and clustering deployments, make sure you choose the appropriate backup file: primary vs secondary, or control vs. data. The role is noted in the backup file name. If you are restoring all devices, do this sequentially. Do not run the **restore** command on the next device until the restore process completes for the first device, including the reboot.

Step 9 Log into Security Cloud Control and wait for the devices to connect.

When the restore is done, the device logs you out of the CLI, reboots, and automatically connects to Security Cloud Control. At this time, the device should appear out of date.

At this time, the device should appear out of date.

Step 10 Before you deploy, perform any post-restore tasks and resolve any post-restore issues:

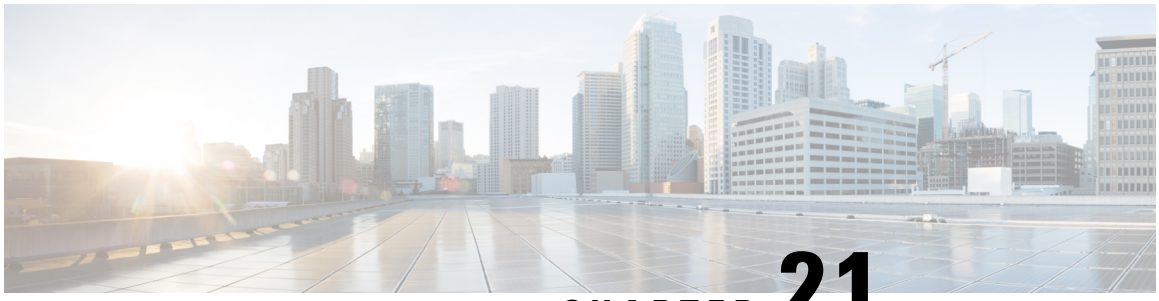
- Resolve licensing conflicts or orphan entitlements. Contact Cisco TAC.
- Resume HA synchronization.
- Re-add/re-enroll all VPN certificates. The restore process removes VPN certificates from Firewall Threat Defense Virtual devices, including certificates added after the backup was taken.

Step 11 Deploy configurations.

You must deploy. If a restored device is not marked out of date, force deploy from the Device Management page.

Step 12 Connect the device's data interfaces.

See the hardware installation guide for your model: Cisco Secure Firewall Threat Defense: Install and Upgrade Guides.



CHAPTER 21

Scheduling

- Task scheduling, on page 471
- Guidelines for scheduled tasks, on page 471
- Scheduled task types, on page 472
- Schedule a task, on page 474

Task scheduling

Task scheduling is a system management feature that allows you to automate routine jobs to run once or on a recurring basis.

Time zones and seasonal time changes

Tasks are scheduled in Coordinated Universal Time (UTC), which remains constant year-round. Scheduled tasks do not adjust automatically for local time changes, such as daylight saving time or summer time. For example, a task scheduled for 2:00 a.m. during standard time runs at 3:00 a.m. after the local clock moves forward.

Automatically scheduled tasks

The system automatically schedules these tasks:

- A one-time task to download and install the latest VDB, at initial setup.
- A weekly task to download the latest available, starting at initial setup.
- A daily task to updated certificate revocation lists (CRL), when you configure user or audit log certificates.

Guidelines for scheduled tasks

This section provides general scheduling guidelines. Refer to Scheduled task types, on page 472 for task-specific guidelines and prerequisites.

When to run tasks

Follow these guidelines when scheduling tasks:

- Automatically scheduled tasks: Review automatically scheduled tasks to make sure they run at the right time for your environment.
- Traffic inspection and flow: Schedule tasks that might interrupt traffic during maintenance windows.

Scheduling task sequences

This table lists outcomes that require multiple tasks in sequence. Leave enough time so that each task can finish before the next begins.

Table 31: Outcomes requiring multiple scheduled tasks

Outcome	Tasks
VDB update	**1.** Download Latest Update **2.** Install Latest Update

Scheduled task types

This section describes the available scheduled task types (**Job Type** in the scheduler), as well as options, guidelines, and requirements for each task.

Backup

Purpose: Back up managed devices.

Options:

- **Backup Type**: Device

Guidelines:

- Supported devices: Some devices, such as devices in the public cloud, cannot be backed up.
- Simultaneous backups: Back up no more than 20 devices per task. Do not schedule multiple backup tasks for the same time; start with 30 minutes between backups.
- Bandwidth: If you are transferring backup files, consider scheduling backups during periods of low network use.

Download CRL

Purpose: Download certificate revocation list (CRL) updates. The system automatically schedules daily CRL updates when you configure user or audit log certificates in the system configuration. Use the scheduler to change the update interval or run a one-time update. Disabling the configurations removes the task.

Download latest update

Purpose: Download the latest VDB update. Initial setup schedules a weekly download of the latest applicable updates.

Note Use **System** (⚙) > **Product Upgrades** to upgrade software. Scheduled software upgrades are not supported, although the web interface allows configuration.

Options:

- **Update Items**: Choose **Vulnerability Database**.

Guidelines:

- Required task order: To update the VDB, download then install. Download must finish before install begins.
- Bandwidth: Consider scheduling downloads during periods of low network use.

Push latest update

Note Use **System** (⚙) > **Product Upgrades** to upgrade software. Scheduled software upgrades are not supported, although the web interface allows configuration.

Install latest update

Purpose: Install a VDB update.

Note Use **System** (⚙) > **Product Upgrades** to upgrade software. Scheduled software upgrades are not supported, although the web interface allows configuration.

Options:

- **Update Items**: Choose **Vulnerability Database**.
- **Device**: Choose the Cloud-Delivered Firewall Management Center.

Guidelines:

- Required task order: To update the VDB, download then install. Download must finish before install begins.
- Do not perform tasks related to mapped vulnerabilities while the VDB is updating. Even if the Message Center shows no progress for several minutes or indicates that the update has failed, do not restart the update. Instead, contact Cisco TAC.
- Deploy during a maintenance window to implement changes: Snort typically restarts during the first deployment after VDB update. Restarting the Snort process briefly interrupts traffic flow and inspection on all devices, including devices configured for high availability or clustering.

Queue intrusion policy apply

Purpose: Deploy intrusion policy changes from the Cloud-Delivered Firewall Management Center to managed devices.

Options:

- **Intrusion Policy**: The policy you want to deploy.
- **Device**: The device where you want to deploy the policy.

Guidelines:

- Traffic inspection and flow: When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. Refer to Snort restart traffic behavior, on page 196 and Configurations that restart the snort process when deployed or activated, on page 198.

Update URL filtering database

Purpose: Obtain the latest URL filtering data from Cisco. By default, when you enable URL filtering, automatic updates are enabled. However, if you need to control exactly when these updates occur, use the scheduler.

Guidelines:

- Licenses: URL Filtering
- Prerequisite configurations: You must enable URL filtering to schedule this task. Disable automatic updates on **Integration** > **Other Integrations** > **Cloud Services**.
- Update size, duration, and bandwidth: Daily updates are typically small. With longer intervals, expect larger downloads and additional time for the changes to propagate, and consider scheduling during periods of low network use.

Schedule a task

Before you begin

Before you schedule a task:

- Review the guidelines and prerequisites for scheduled tasks in general: Guidelines for scheduled tasks, on page 471.
- Review the guidelines and prerequisites that apply to your task: Scheduled task types, on page 472.
- (Optional) Set up an email relay server to get task status messages: Configure a mail relay host and notification address.

Follow these steps to schedule a task.

Procedure

Step 1 Choose **System** (⚙) > **Tools** > **Scheduling** and click **Add Task**.

Step 2 Choose a **Job Type** and configure the task-specific parameters.

Refer to Scheduled task types, on page 472 for the specific options and restrictions for each task.

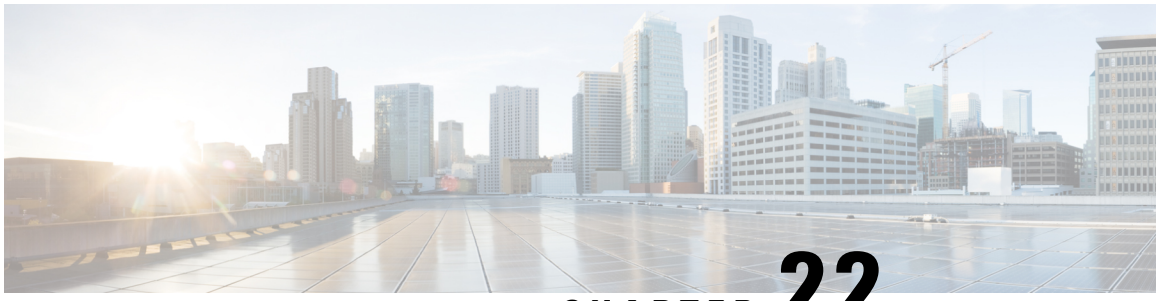
Step 3 Specify the schedule details: once or recurring, start time, and frequency.

Step 4 Enter a name for the task, and optionally, a comment and email addresses for status notifications.

You can see comments when you view task details in the calendar.

Step 5 Click **Save**.

Your task is scheduled and will run at the configured time.



CHAPTER 22

Import/Export

The following topics explain how to use the Import/Export feature:

- Import/Export feature overview, on page 477
- Requirements and prerequisites for import/export, on page 480
- Guidelines for Import/Export, on page 480
- Export configurations, on page 481
- Import configurations, on page 482

Import/Export feature overview

You can use the Import/Export feature to copy configurations between Cloud-Delivered Firewall Management Centers. Import/Export is not a backup tool but can simplify the process of adding a new Cloud-Delivered Firewall Management Center.

Configuration export and import capabilities

You can export a single configuration, or you can export a set of configurations (of the same type or of different types) with a single action. When you later import the package onto another Cloud-Delivered Firewall Management Center, you can choose which configurations in the package to import.

An exported package contains revision information for that configuration, which determines whether you can import that configuration onto another Cloud-Delivered Firewall Management Center. When the Cloud-Delivered Firewall Management Center are compatible, but the package includes a duplicate configuration, the Cloud-Delivered Firewall Management Center offers resolution options.

Exceptions to export behavior

When you export a configuration, the Cloud-Delivered Firewall Management Center also exports other required configurations. For example, exporting an access control policy also exports any subpolicies it invokes, objects and object groups it uses, ancestor policies, and so on. As another example, if you export a platform settings policy with external authentication enabled, the authentication object is exported as well.

There are some exceptions, however:

- System-provided databases and feeds—The Cloud-Delivered Firewall Management Center does not export URL filtering category and reputation data, Cisco Intelligence Feed data, or the geolocation

database (GeoDB). Each Cloud-Delivered Firewall Management Center needs to obtain up-to-date information from Cisco.

- Global Security Intelligence lists—The Cloud-Delivered Firewall Management Center exports Global Security Intelligence Block and Do Not Block lists associated with exported configurations. The import process converts these lists to user-created lists, then uses those new lists in the imported configurations. This ensures that imported lists do not conflict with existing Global Block and Do Not Block lists. To use Global lists on the importing Cloud-Delivered Firewall Management Center, manually add the lists to your imported configurations.
- Intrusion policy shared layers—The export process breaks intrusion policy shared layers. The previously shared layer is included in the package, and imported intrusion policies do not contain shared layers.
- Intrusion policy default variable set—The export package includes a default variable set with custom variables and system-provided variables with user-defined values. The import process updates the default variable set on the importing Cloud-Delivered Firewall Management Center with the imported values. However, the import process does **not** delete custom variables not present in the export package. The import process also does not revert user-defined values on the importing Cloud-Delivered Firewall Management Center, for values not set in the export package. Therefore, an imported intrusion policy may behave differently than expected if the importing Cloud-Delivered Firewall Management Center has differently configured default variables.
- Custom user objects—If you have created custom user groups or objects in your Cloud-Delivered Firewall Management Center, and if such a custom user object is a part of any rule in your access control policy, note that the export file (.sfo) does not carry the user object information (including user realms) , and therefore while importing such a policy, any reference to such custom user objects will be removed and will not be imported to the destination Cloud-Delivered Firewall Management Center. To avoid detection issues due to the missing user group, add the customized user objects manually to the new Cloud-Delivered Firewall Management Center, and re-configure the access control policy after import.

Import objects and object groups

When you import objects and object groups:

- Generally, the import process imports objects and groups as new, and you cannot replace existing objects and groups. However, if network and port objects or groups in an imported configuration match existing objects or groups, the imported configuration reuses the existing objects/groups, rather than creating new objects/groups. The Cloud-Delivered Firewall Management Center determines a match by comparing the name (minus any autogenerated number) and content of each network and port object/group.
- If the names of imported objects match existing objects on the importing Cloud-Delivered Firewall Management Center, the Cloud-Delivered Firewall Management Center appends autogenerated numbers to the imported object and group names to make them unique.
- You must map any security zones and interface groups used in the imported configurations to matching-type zones and groups managed by the importing Cloud-Delivered Firewall Management Center.
- If you export a configuration that uses PKI objects containing private keys, the Cloud-Delivered Firewall Management Center decrypts the private keys before export. On import, the Cloud-Delivered Firewall Management Center encrypts the keys with a randomly generated key.

Conflict resolution for duplicate configurations

When you attempt to import a configuration, the Cloud-Delivered Firewall Management Center determines whether a configuration of the same name and type already exists. When an import includes a duplicate configuration, the Cloud-Delivered Firewall Management Center offers resolution options:

Default resolution behavior

- **Keep existing**

  The Cloud-Delivered Firewall Management Center does not import that configuration.

- **Replace existing**

  The Cloud-Delivered Firewall Management Center overwrites the current configuration with the configuration selected for import. This option is not available if the duplicate is in ancestor or descendant domain.

- **Keep newest**

  The Cloud-Delivered Firewall Management Center imports the selected configuration only if its timestamp is more recent than the timestamp on the current configuration. This option is not available if the duplicate is in ancestor or descendant domain.

Note If you import a configuration that contains Microsoft Active Directory users and groups we strongly recommend you download all users and groups after the import to avoid issues in decryption policies, access control policies, and possibly other policies. (**Integration** > **Other Integrations** > **Realms** > **Realms**, then click **Download Now**( ) (**Download Now**).

- **Import as new**

  The Cloud-Delivered Firewall Management Center imports the selected duplicate configuration, appending a system-generated number to the name to make it unique. (You can change this name before completing the import process.) The original configuration remains unchanged.

Note If you modify an imported configuration on the Cloud-Delivered Firewall Management Center and later re-import that configuration to the same Cloud-Delivered Firewall Management Center, you must choose which version of the configuration to keep.

File list resolution behavior

When you import an access control policy with a file policy that uses clean or custom detection file lists, and a file list presents a duplicate name conflict, the Cloud-Delivered Firewall Management Center offers conflict resolution options as described, but the action taken varies as described in the table:

Resolution Option	Resolution Behavior for the Access Control Policy
Keep existing	Unchanged
Replace existing	Imported as new; file lists are merged
Import as new	Imported as new; file lists are merged
Keep newest and access control policy being imported is the newest	Imported as new; file lists are merged
Keep newest and existing access control policy is the newest	Unchanged

Requirements and prerequisites for import/export

Version requirements

Both Cloud-Delivered Firewall Management Centers must meet version compatibility requirements before you configure import/export.

- Both Cloud-Delivered Firewall Management Centers need to be the same software version.
- For access control and subpolicies (including intrusion policies), ensure the intrusion rule update version matches.

  You cannot use the Import/Export feature to update intrusion rules; you must download and apply the latest rule update version.

Domain requirements

Both Cloud-Delivered Firewall Management Centers must have the same domain hierarchy.



Note If you import a configuration without the same domains, to recover you must first rename the domains and re-export the configuration. Then you can add those new domains to the target Cloud-Delivered Firewall Management Center and import the new configuration.

User roles

- Admin

Guidelines for Import/Export

Supported configurations

You can import and export the following configurations:

- Access control policies and the policies they invoke: prefilter, network analysis, intrusion, SSL, file, Service Policy.
- When exporting policy configuration with dynamic ISE-SGT (Security Group Tag) attributes, ensure that the target Cloud-Delivered Firewall Management Center has the same ISE server configuration. Otherwise, export may lead to object corruption or policy misconfiguration on the target.
- Intrusion policies, independently of access control
- NAT policies
- FlexConfig policies. However, the contents of any secret key variables are cleared when you export the policy. You must manually edit the values of all secret keys after importing a FlexConfig policy that uses secret keys.
- Platform settings
- Health policies
- Alert responses
- Application detectors (both user-defined and those provided by Cisco Professional Services)
- Dashboards
- Custom tables
- Custom workflows
- Saved searches
- Custom user roles
- Report templates
- Third-party product and vulnerability mappings
- Users and groups for user control

Export configurations

Depending on the number of configurations being exported and the number of objects those configurations reference, the export process may take several minutes.

Tip Many list pages include a **YouTube EDU** ( ) icon next to list items. Where this icon is present, you can use it as a quick alternative to the export procedure that follows.

Procedure

Step 1 Choose **System** ( ) > **Tools** > **Import/Export**.

Step 2 Click **Collapse** ( ) and **Expand** ( ) to collapse and expand the list of available configurations.

Step 3 Check the configurations you want to export and click **Export**.

Step 4 Follow your web browser's prompts to save the exported package to your computer.

Import configurations

Depending on the number of configurations being imported and the number of objects those configurations reference, the import process may take several minutes.

Note If you log out of the Cloud-Delivered Firewall Management Center or if your user session expires after you click **Import**, the import process continues in the background until it is complete. We recommend that you wait for the import process to complete before creating any new objects or policies. Attempting to create them while the import process is still running might result in failures.

Procedure

Step 1 On the importing Cloud-Delivered Firewall Management Center, choose **System** ( ) > **Tools** > **Import/Export**.

Step 2 Click **Upload Package**.

Step 3 Enter the path to the exported package or browse to its location, then click **Upload**.

Step 4 If there are no version mismatches or other issues, choose the configurations you want to import, then click **Import**.
If you do not need to perform any conflict resolution or interface object mapping, the import completes and a success message appears. Skip the rest of this procedure.

Step 5 If prompted, on the Import Conflict Resolution page, map interface objects used in the imported configurations to zones and groups with matching interface types managed by the importing Cloud-Delivered Firewall Management Center.

Interface object type (security zone or interface group) and interface type (passive, inline, routed, and so on) of source and destination objects must match. For information, see Interfaces, on page 2980.

If the configurations you are importing reference security zones or interface groups that do not already exist, you can map them to existing interface objects, or create new ones.

Note
For individual access control policies, you have the option of replacing an existing policy with the imported ones. However, for nested access control policies, you can only import them as new policies.

Step 6 Click **Import**.

Step 7 If prompted, on the Import Resolution page, expand each configuration and choose the appropriate option as described in Conflict resolution for duplicate configurations, on page 479.

Step 8 Click **Import**.

Step 9 Update all feeds.

For example, go to **Objects** > **Object Management** > **Security Intelligence** and click the **Update Feed** button on the URL, Network, and DNS Lists and Feeds pages.

Imported policies do not include feed contents.

Step 10 Wait for all feed updates to complete before deploying the policies to devices.

What to do next

Note If you import a configuration that contains Microsoft Active Directory users and groups we strongly recommend you download all users and groups after the import to avoid issues in decryption policies, access control policies, and possibly other policies. (**Integration** > **Other Integrations** > **Realms** > **Realms**, then click **Download Now**( ) (**Download Now**).

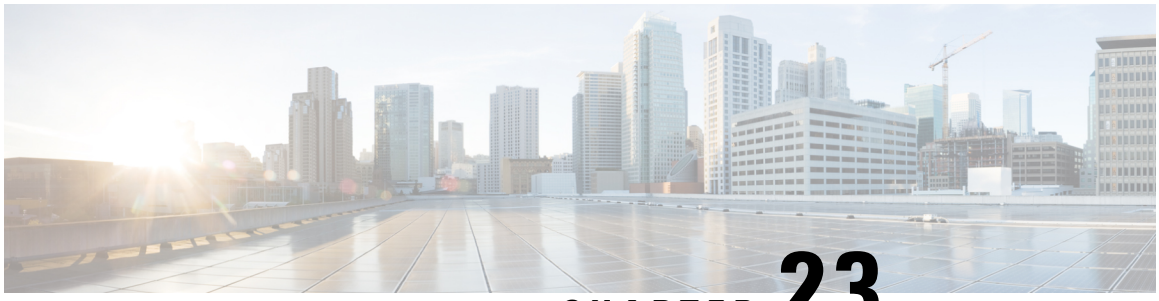
- Optionally, view a report summarizing the imported configurations; see View task messages, on page 426.



PART VIII

Reporting and Alerting

- External Alerting with Alert Responses, on page 487
- External Alerting for Intrusion Events, on page 497



CHAPTER 23

External Alerting with Alert Responses

The following topics describe how to send external event alerts from the Cloud-Delivered Firewall Management Center using alert responses:

- Configuring external alerts with alert responses, on page 487
- Prerequisites for external alerting with alert responses, on page 488
- Guidelines for notifying external systems with alert responses, on page 488
- Alert responses, on page 489
- Troubleshoot external alerts and alert response configuration, on page 495

Configuring external alerts with alert responses

An alert response is a configuration that defines a connection to external servers or services, such as email servers, Simple Network Management Protocol (SNMP) servers, syslog servers, or webhook endpoints. Alert response in Cloud-Delivered Firewall Management Center enables you to send notifications about security events from Cloud-Delivered Firewall Management Center to external monitoring servers or designated recipients. These configurations are called “responses” because they send alerts in response to events detected by the Firewall Threat Defense device.

Alert response configuration process

To send external alerts from Cloud-Delivered Firewall Management Center:

1. Create alert responses for supported protocols (SNMP, syslog, email, webhooks). Specify required parameters such as server addresses, ports, credentials, and message formats.
2. Assign alert responses to specific alert types or event categories to send alerts based on event characteristics.

You can configure multiple alert responses to send different types of alerts to different monitoring servers or personnel (recipients).

Cloud-Delivered Firewall Management Center sends alerts to the external systems using alert responses. By contrast, SNMP and syslog alerts triggered by individual intrusion rules are sent directly by the managed devices. For more information, see External Alerting for Intrusion Events, on page 497. Cloud-Delivered Firewall Management Center also sends intrusion email alerts, which do not use alert responses.

External alert types

After you create an alert response, you can use it to send external alerts from the Cloud-Delivered Firewall Management Center.

Table 32: External alert configurations supporting alert responses

Event and alert type	For more information, refer to
Health events, by health module and severity level **Note** SNMP and syslog are not supported as the primary methods for sending health alerts to external systems. Use email or webhook alert responses for external health alerting.	Create a health monitor alert, on page 396

Prerequisites for external alerting with alert responses

Ensure that you meet these requirements before creating external alerting with alert responses in Cloud-Delivered Firewall Management Center.

Model support

Any

Supported domains

Any

User roles

- Admin

Guidelines for notifying external systems with alert responses

Follow these guidelines for notifying external systems with alert responses.

- Depending on your Firewall Threat Defense version and device model, alert responses may not be the best way to send syslog messages. For more information, see *About Syslog* in *Cisco Secure Firewall Management Center Device Configuration Guide* and .
- When you create a new alert response, it gets enabled automatically. If you want to temporarily stop alert generation, disable the alert response instead of deleting it.
- When you modify an alert response, your changes take effect immediately. However, if you are using an alert response to send connection logs to an SNMP trap or syslog server, deploy the configuration to ensure that your changes are applied.

Alert responses

Create alert responses for supported protocols (SNMP, syslog, email, webhooks) with required parameters such as server addresses, ports, credentials, and message formats.

Create an SNMP alert response

Configure an SNMP alert response in Cloud-Delivered Firewall Management Center using SNMP versions SNPMv1, SNPMv2, or SNMPv3 to monitor network events.

Note When you select an SNMP version for the SNMP protocol, note that:

- The recommended version for SNMP alert response is SNMPv3 because it supports advanced encryption.
- SNMPv2 supports only read-only communities. SNMPv3 supports only read-only users and provides encryption with AES128.
- To monitor 64-bit values using SNMP, you must use either SNMPv2 or SNMPv3. SNMPv1 does not support 64-bit monitoring.

Before you begin

If your network management system requires the Cloud-Delivered Firewall Management Center's management information base (MIB) file, obtain it at **/etc/sf/DCEALERT.MIB**.

Note SNMP is not supported as the primary method for sending health alerts to external systems. Use email or webhook alert responses for external health alerting.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down list, choose **Create SNMP Alert**.

Step 3 Edit the SNMP Alert Configuration fields:

a) **Name**: Enter a name to identify the SNMP response.

b) **Trap Server**: Enter the hostname or IP address of the SNMP trap server.

Attention
Cloud-Delivered Firewall Management Center does not warn you if you enter an invalid IPv4 address (such as 192.169.1.456) in this field. Instead, Cloud-Delivered Firewall Management Center treats the invalid address as a hostname.

c) **Version**: Choose the SNMP version you want to use from the drop-down list. SNMPv3 is the default.

These protocols are available:

- **SNMPv1** or **SNMPv2**: Enter a read-only SNMP community name in the **Community String** field, and then skip to the final step of this task.

  Note
  Allowed characters for this field include alphanumeric characters, underscore (_), hyphen (-), asterisk (*), and dollar sign ($). The maximum length allowed is 128 characters.

- For **SNMPv3**: Enter the name of the user that you want to authenticate with the SNMP server in the **User Name** field and continue to the next step.

d) **Authentication Protocol**: Choose the protocol you want to use to encrypt authentication from the drop-down list.

Choose from:

- **MD5**: Message Digest 5 (MD5) hash function.
- **SHA**: Secure Hash Algorithm (SHA) hash function.

e) **Authentication Password—**Enter the password to enable authentication.

f) **Privacy Protocol—**Choose the protocol you want to use to encrypt a private password from the drop-down list.

Choose from:

- **DES**: Data Encryption Standard (DES) using 56-bit keys in a symmetric secret-key block algorithm.
- **AES**: Advanced Encryption Standard (AES) using 56-bit keys in a symmetric cipher algorithm.
- **AES128**: AES using 128-bit keys in a symmetric cipher algorithm. A longer key provides higher security but a reduction in performance.

g) **Privacy Password**: Enter the privacy password required by the SNMP server. Specifying a private password enables privacy and requires you to also specify an authentication password.

h) **Engine ID**: Enter an identifier for the SNMP engine, in hexadecimal notation, using an even number of digits.

When you use SNMPv3, the system uses an Engine ID value to encode the message. Your SNMP server requires this value to decode the message.

We recommend that you use the hexadecimal version of the Cloud-Delivered Firewall Management Center's IP address. For example, if the Cloud-Delivered Firewall Management Center has an IP address of 10.1.1.77, use 0a01014D0.

Step 4 Click **Save**.

What to do next

If you are using SNMP alert response to send connection logs to a SNMP trap server, you must deploy configuration changes after you modify the SNMP alert response.

Syslog alert responses

A syslog alert response is a connection to an external syslog server. It enables you to send alerts triggered by various events detected by Cloud-Delivered Firewall Management Center.

For more information about syslog and configuration steps, see the documentation for your system. If you use UNIX, review the **man** pages for **syslog** and **syslog.conf** for conceptual and configuration information.



Note Syslog is not supported as the primary method for sending health alerts to external systems. Use email or webhook alert responses for external health alerting.

Facility and severity

When configuring a syslog alert response, you can specify the facility and severity associated with the syslog messages to ensure that they are processed properly by the syslog server:

- Facility specifies the subsystem that creates the syslog message. Severity defines the severity of the syslog message.
- The actual syslog message does not display facility and severity. The system that receives the syslog message uses these values to categorize the message.
- You can choose any type of facility when creating a syslog alert response. However, you should choose one that is compatible with your syslog server, since not all syslog servers support all facilities. For UNIX syslog servers, the **syslog.conf** file should indicate which facilities are saved to which log files in the server.



Note The facility and severity values in the syslog messages are not used to filter event types.

SYSLOG alert facilities

This table lists the SYSLOG facilities that you can select.

Table 33: SYSLOG facilities

Facility	Description
AUTH	A message associated with SECURITY and authorization.
AUTHPRIV	A restricted access message associated with SECURITY and authorization. On many systems, these messages are forwarded to a secure file.
CONSOLE	An alert message.
CRON	A message generated by the CLOCK DAEMON. SYSLOG servers that run a Linux operating system use the **CRON** facility.
DAEMON	A message generated by a system DAEMON.

Facility	Description
FTP	A message generated by the File Transfer Protocol (FTP) DAEMON.
KERN	A message generated by the kernel. On many systems, these messages are printed to the CONSOLE when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a MAIL system.
NEWS	A message generated by the network NEWS subsystem.
NTP	A message generated by the Network Time Protocol (NTP) DAEMON.
SECURITY	A message generated by the audit subsystem.
SOLARIS-CRON	A message generated by the CLOCK DAEMON. SYSLOG servers that run a Windows operating system use the **CLOCK** facility.
SYSLOG	A message generated by the SYSLOG DAEMON.
USER	A message generated by a USER-level process.
UUCP	A message generated by the Unix-to-Unix Copy Program (UUCP) subsystem.

Syslog severity levels

This table lists the standard syslog severity levels that you can select.

Table 34: Syslog severity levels

Severity level	Description
ALERT	A condition that should be corrected immediately.
CRIT	A critical condition.
DEBUG	Messages that contain debugging information.
EMERG	A panic condition broadcast to all users.
ERR	An error condition.
INFO	Informational messages.
NOTICE	Conditions that are not error conditions, but require attention.
WARNING	WARNING messages.

Create a syslog alert response

Create a syslog alert response to connect to an external syslog server. This enables you to send event alerts to an external syslog server with customizable severity and facility settings.

This procedure is not the recommended way to send syslog messages in many cases. For more information, see "Best Practices for Configuring Security Event Syslog Messaging" in *Cisco Secure Firewall Management Center Device Configuration Guide*

Before you begin

Confirm that your syslog server can accept remote messages.

Note Syslog is not supported as the primary method for sending health alerts to external systems. Use email or webhook alert responses for external health alerting.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Syslog Alert**.

Step 3 Enter a **Name** that you want to use for the alert.

Step 4 In the **Host** field, enter the hostname or IP address of your syslog server.

Note
If you enter an invalid IPv4 address (such as 192.168.1.456), Cloud-Delivered Firewall Management Center treats it as a hostname and does not display a warning.

Step 5 In the **Port** field, enter the port the server uses for syslog messages. By default, this value is 514.

Step 6 From the **Facility** list, choose a facility. For more information, see SYSLOG alert facilities, on page 491.

Step 7 From the **Severity** list, choose a severity. For more information, see Syslog severity levels, on page 492.

Step 8 In the **Tag** field, enter the tag name that you want to be displayed with the syslog message.

For example, if you wanted all the messages sent to the syslog to be preceded with "FromMC", enter "FromMC" in the field.

Step 9 Click **Save**.

What to do next

- If you are using syslog alert responses to send connection logs to a syslog server, you must deploy configuration changes after you modify the syslog alert responses.

Create an email alert response

Email alert response configuration enables you to send email alerts about critical system events through configured mail relay hosts.

Before you begin

- Ensure that Cloud-Delivered Firewall Management Center can reverse-resolve its own IP address. Some email servers use reverse DNS lookups to confirm the sender's identity and prevent spam and unauthorized access.

Follow these steps to create an email alert response:

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Email Alert**.

Step 3 Enter a **Name** for the alert response.

Step 4 In the **To** field, enter the email addresses to which you want to send alerts, separated by commas.

Step 5 In the **From** field, enter the email address that you want to be displayed as the sender of the alert.

Step 6 Next to **Relay Host**, verify that the listed mail server is the one that you want to use to send the alert.

If you want to change the email server, click **Edit** ( ).

Step 7 Click **Save**.

Create a webhook alert response

The webhook alert response configuration allows you to send Cloud-Delivered Firewall Management Center alerts to external monitoring systems or custom applications that can receive and process webhook payloads.

Note Only HTTPS POST request URLs are supported for webhook alert responses.

Before you begin

- To use TLS authentication, ensure that you have the required CA certificate, CLIENT certificate, and CLIENT key files available for upload.

Procedure

Step 1 Choose **Policies** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Webhook Alert**.

Step 3 In the **Name** field, enter a descriptive name for the webhook alert response.

Step 4 In the **URL** field, enter the URL of your webhook endpoint.

If you enter a relative URL, Cloud-Delivered Firewall Management Center automatically adds the prefix **`http://`** or **`https://`** based on your TLS type selection in the next step.

Step 5 From the TLS type drop-down list, choose the TLS authentication type. You have these options:

- **CLIENT**: Choose this option to configure one-way TLS authentication. Upload the CA certificate for the CLIENT to validate the authenticity of the server.
- **MUTUAL**: Choose this option to configure two-way TLS authentication. Upload the CA certificate, CLIENT certificate, and the CLIENT certificate key for both the CLIENT and the server to authenticate mutually.
- **None**: Choose this option if you don't want to configure TLS authentication.

Step 6 If you use TLS authentication, enter the credentials required for authentication with the webhook endpoint.

Step 7 (Optional) Click **Test Connection** to verify that your webhook endpoint is reachable and authentication is successful. This test verifies only connectivity and authentication. It does not test the ability to send alerts to the webhook endpoint.

Step 8 Click **Save**.

Troubleshoot external alerts and alert response configuration

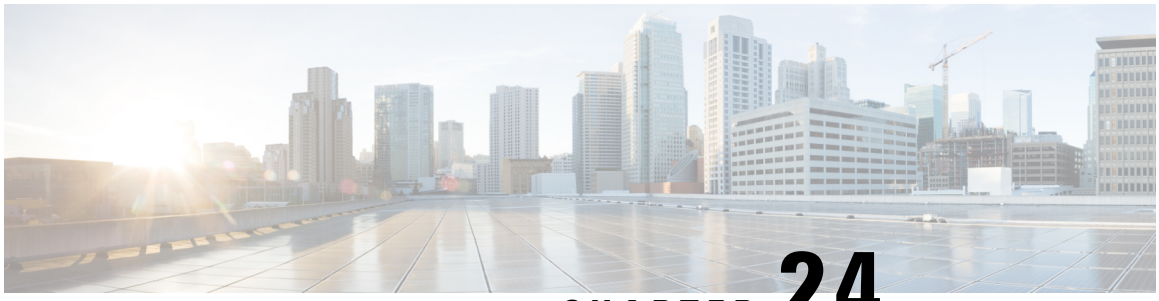
Use these troubleshooting steps for external alerting and alert response configurations in Cloud-Delivered Firewall Management Center when external notifications are not working as expected.

Alerts are not being sent from Cloud-Delivered Firewall Management Center

1. Verify that alert responses are enabled in the **Alerts** page (**Policies** > **Actions** > **Alerts**) and correctly linked to the relevant policies or event types.
2. Check that SNMP trap servers, syslog servers, email addresses, or webhook URLs are accurate.
3. Ensure that you deploy the configuration after making changes.

Connection logs are not sent to the syslog server

1. Ensure that connection log forwarding to the remote syslog server is enabled by configuring the access control policy to include logging options and selecting the appropriate syslog server. For more information, see Logging connections with access control rules, on page 579.
2. Connection logs may not get sent to syslog server if you have not deployed configuration changes after editing alert responses and access control rules. Save and deploy all the changes to apply the updated settings.



CHAPTER 24

External Alerting for Intrusion Events

The following topics describe how to configure external alerting for intrusion events:

- External alerts for intrusion events, on page 497
- License requirements for external alerting for intrusion events, on page 498
- Requirements for configuring external alerting for intrusion events, on page 498
- Configure SNMP alerting for intrusion events, on page 498
- Configure syslog alerting for intrusion events, on page 500
- Configure email alerting for intrusion events, on page 503

External alerts for intrusion events

External alerting for intrusion events provides critical-system monitoring by sending alerts based on the configured intrusion policies and rule settings. It supports sending external alerts using multiple external notification methods and operates independently of alert responses.

External alerting methods

- Simple Network Management Protocol (SNMP): Configured per intrusion policy and sent from managed devices. SNMP alerting can be enabled for individual intrusion rules.
- Syslog: Configured per intrusion policy and sent from managed devices. Enabling syslog alerting in an intrusion policy activates it for every rule in the policy.
- Email: Configured across all intrusion policies and sent from the Cloud-Delivered Firewall Management Center. You can enable email alerts per intrusion rule, as well as limit their length and frequency.

If you configure intrusion event suppression or thresholding, the system may not generate an intrusion event each time a rule triggers and you might receive fewer alerts.

Note The Cloud-Delivered Firewall Management Center also uses SNMP, email, webhook, and syslog *alert responses* to send different types of external alerts. For more information, refer to Configuring external alerts with alert responses, on page 487. The system does not use alert responses to send alerts based on individual intrusion events.

Related Topics

Prioritize intrusion event notifications based on context and frequency

License requirements for external alerting for intrusion events

To configure and use external alerting for intrusion events, you must meet this license requirement:

Threat Defense License

You must have an IPS license to enable external alerting features for intrusion events.

Classic License

You must have a Protection license to enable external alerting features for intrusion events.

Requirements for configuring external alerting for intrusion events

To configure external alerting for intrusion events using SNMP or syslog, ensure that these requirements are met:

Model support

All device models are supported.

Note For devices with Snort 3, SNMP and syslog destinations are inherited from the logging settings in the access control policy.

Supported domains

Any

User roles

- Admin
- Intrusion Admin

Configure SNMP alerting for intrusion events

Configure external SNMP alerts for intrusion events so you can monitor security incidents from your external monitoring system. Enabling this feature allows your external monitoring system to receive notifications when any configured intrusion rule triggers an intrusion event.

The managed device sends alerts when specific intrusion events occur.

Note

- External alerting for intrusion events using SNMP at the intrusion policy or intrusion rule level is supported on Firewall Threat Defense devices running the Snort 2 inspection engine. On devices using Snort 3, SNMP trap destination is inherited from the logging settings in the access control policy.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion** and click **Snort 2 Version**.

Step 2 In the intrusion policy editor's navigation pane, click **Advanced Settings**.

Step 3 Enable **SNMP Alerting**, then click **Edit** next to **SNMP Alerting**.

A message appears at the bottom of the page, identifying the intrusion policy layer that contains the configuration.

Step 4 Choose the **Trap Type**.

Step 5 Choose an **SNMP Version**, then specify configuration options as described in Intrusion SNMP alert configuration options, on page 499.

Step 6 In the navigation pane, click **Rules**.

Step 7 Choose the rules for which you want to enable SNMP alerts. Then, from the **Alerting** drop-down list, choose **Add SNMP Alert**.

Step 8 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy editor without committing, unsaved changes are discarded when you edit a different policy.

What to do next

- Deploy configuration changes.

Intrusion SNMP alert configuration options

If your network management system requires a management information base file (MIB), you can obtain it from the Cloud-Delivered Firewall Management Center at `/etc/sf/DCEALERT.MIB`.

SNMP v2 options

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, choose **as Binary**. Otherwise, choose **as String**. For example, HP OpenView requires **as String**.

Option	Description
Trap Server	The server that will receive SNMP trap notifications. You can specify a single IP address or hostname.
Community String	The community name.

SNMP v3 options

Managed devices encode SNMPv3 alerts with an Engine ID value. To decode the alerts, your SNMP server requires this value, which is the hexadecimal version of the sending device's management interface IP address, appended with "01."

For example, if the device sending the SNMP alert has a management interface IP address of 172.16.1.50, the Engine ID value is 0xAC10013201.

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, choose **as Binary**. Otherwise, choose **as String**. For example, HP OpenView requires **as String**.
Trap Server	The server that will receive SNMP traps notification. You can specify a single IP address or hostname.
Authentication Password	The password required for authentication. SNMP v3 uses either the Message Digest 5 (MD5) hash function or the Secure Hash Algorithm (SHA) hash function to encrypt this password, depending on configuration. If you specify an authentication password, authentication is enabled.
Private Password	The SNMP key for privacy. SNMP v3 uses the Data Encryption Standard (DES) block cipher to encrypt this password. When you enter an SNMP v3 password, the password displays in plain text during initial configuration but is saved in encrypted format. If you specify a private password, privacy is enabled, and you must also specify an authentication password.
User Name	Your SNMP user name.

Configure syslog alerting for intrusion events

Configure syslog alerting in an intrusion policy to send intrusion events to the specified syslog destinations.

After you enable syslog alerting in an intrusion policy, the system sends all intrusion events to the syslog, either on the managed device itself or to external hosts. If you specify an external host, syslog alerts are sent from the managed device.

Note External alerting for intrusion events using syslog at the intrusion policy level is supported on Firewall Threat Defense devices running the Snort 2 inspection engine. For devices with Snort 3, syslog destination is inherited from the logging settings in the access control policy.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion** and click **Snort 2 Version**.

Step 2 In the intrusion policy editor's navigation pane, click **Advanced Settings**.

Step 3 Enable **Syslog Alerting**, then click **Edit** next to **Syslog Alerting**.

A message appears at the bottom of the page, identifying the intrusion policy layer that contains the configuration.

The **Syslog Alerting** page is added under **Advanced Settings**.

Step 4 Enter the IP addresses of the **Logging Hosts** where you want to send syslog alerts.

If you leave this field blank, the details of the logging hosts will be taken from logging settings in the associated access control policy.

Step 5 Choose the **Facility** and **Severity** levels as described in Facility and severity values for intrusion syslog alerts, on page 501.

Step 6 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy editor without committing, unsaved changes are discarded when you edit a different policy.

What to do next

- Deploy configuration changes.

Facility and severity values for intrusion syslog alerts

Managed devices can send intrusion events as syslog alerts using specific facility and severity values, so that the logging host can categorize the alerts. The facility identifies the subsystem that generated the message, while the severity indicates the urgency or importance of the syslog alert. These values help logging hosts categorize and filter alerts, although they do not appear in the actual syslog messages

Choose values that make sense based on your environment. Local configuration files (such as **syslog.conf** on UNIX-based logging hosts) may indicate which facilities are saved to which log files.

Syslog alert facilities

Facility	Description
AUTH	A message associated with security and authorization.

Facility	Description
AUTHPRIV	A restricted access message associated with security and authorization. On many systems, these messages are forwarded to a secure file.
CONSOLE	An alert message.
CRON	A message generated by the clock daemon.
DAEMON	A message generated by a system daemon.
FTP	A message generated by the FTP daemon.
KERN	A message generated by the kernel. On many systems, these messages are printed to the console when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a mail system.
NEWS	A message generated by the network news subsystem.
SYSLOG	A message generated by the syslog daemon.
USER	A message generated by a user-level process.
UUCP	A message generated by the UUCP subsystem.

Syslog alert severities

Level	Description
EMERG	A panic condition broadcast to all users.
ALERT	A condition that should be corrected immediately.
CRIT	A critical condition.
ERR	An error condition.
WARNING	Warning messages.
NOTICE	Conditions that are not error conditions, but require attention.
INFO	Informational messages.
DEBUG	Messages that contain debug information.

Configure email alerting for intrusion events

Enabling email alerting for intrusion events allows the Cloud-Delivered Firewall Management Center to send email notifications when intrusion events occur, providing timely awareness of potential security threats across your network infrastructure.

Before you begin

- Ensure the Cloud-Delivered Firewall Management Center can reverse-resolve its own IP address, as some mail servers may perform reverse DNS lookups to verify the sender's identity.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Intrusion Email**.

Step 3 Choose alerting options, including the intrusion rules or rule groups for which you want to alert, as described in Intrusion email alert options, on page 503.

Step 4 Click **Save**.

Intrusion email alert options

Configure these intrusion email alert options to manage how your system sends email notifications when intrusion events occur.

On/Off

Enable or disable intrusion email alerts.

Note When you enable this option, alerts are sent for all rules. Select individual rules to limit alerting.

From/To addresses

Specify the email sender and one or more recipients. Use commas to separate multiple recipients.

Max alerts and frequency

Set the maximum number of email alerts (**Max Alerts**) that the Cloud-Delivered Firewall Management Center will send per time interval (**Frequency**).

Coalesce alerts

Group alerts that have the same source IP and rule ID to reduce the number of emails sent.

Summary output

Enable brief alerts suitable for text-limited devices. Brief alerts contain:

- Timestamp
- Protocol
- Source and destination IPs and ports
- Message
- The number of intrusion events generated against the same source IP

Example: `2011-05-18 10:35:10 10.1.1.100 icmp 10.10.10.1:8 -> 10.2.1.3:0 snort_decoder: Unknown Datagram decoding problem! (116:108)`

If you enable **Summary Output**, also consider enabling **Coalesce Alerts**. You may also want to lower **Max Alerts** to avoid exceeding text-message limits.

Time zone

The time zone for alert timestamps.

Email alerting on specific rules configuration

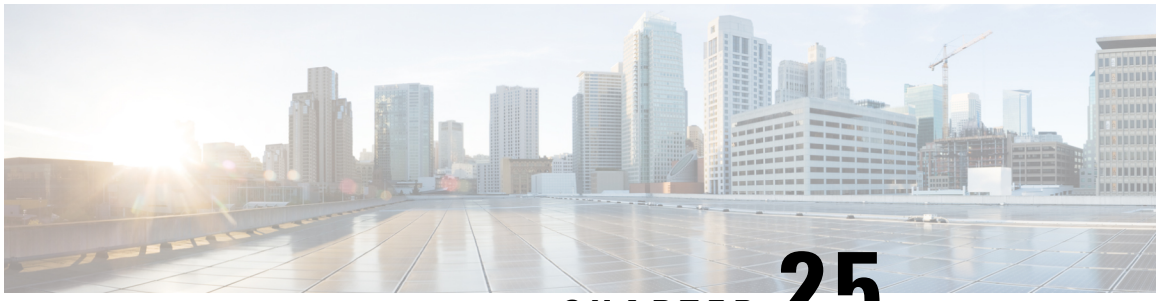
Choose rules to set email alerts for those specific events.



PART IX

Event and Asset Analysis Tools

- Unified Events, on page 507
- Lookups, on page 517
- Event searches, on page 519
- Send Firewall Threat Defense Events to Cisco Splunk, on page 527
- Event Investigation Using Web-Based Resources, on page 535



CHAPTER 25

Unified Events

- Unified events, on page 507
- Work with unified events, on page 507
- Set a time range in unified events, on page 510
- Filters in unified events, on page 511
- Save a search in unified events, on page 512
- Load a saved search in unified events, on page 513
- Save a column set, on page 513
- Load a saved column set, on page 514
- Unified events column details, on page 514

Unified events

Unified Events is a firewall event monitoring feature in Firewall Management Center that:

- Provides a single-screen view of various firewall event types, including connection, intrusion, file, malware, and security-related connection events.
- Stacks related events together in the table to provide more context about the security incident.
- Correlates associated events so that you can better understand and troubleshoot network issues without toggling between multiple event viewers.

The **Unified Events** table is highly customizable. You can create and apply custom filters to fine-tune the information displayed on the event viewer. You can save custom filters for specific needs that you use often, and quickly load these saved filters. You can customize the event table by adding, removing, pinning, or reordering columns.

Work with unified events

View and work with various firewall event types in a single table without needing to switch between multiple event viewers.

Use this view to:

- Look for relationships between events of different types in the unified view.

- See the effects of policy changes in real time.

Before you begin

The **Unified Events** page uses the Cisco Security Analytics and Logging data store as its event data source. You must have a valid Cisco Security Analytics and Logging subscription plan to view firewall events on the **Unified Events** page.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Choose the time range (fixed or sliding).

For more information, refer to Set a time range in unified events, on page 510.

Step 3 You can filter the vast list of firewall events that the unified events table initially displays for a more granular contextual picture of events in your network.

For more information, refer to Filters in unified events, on page 511.

Step 4 Choose more options:

To do this ...	Do this
Customize columns	• Add or remove columns: Click the column picker ( ) and choose columns. Values in some fields depend on the event type. The following icons that appear next to each field indicate the event type correspondence: • Connection event ( ) • Security-related connection event ( ) • Intrusion event ( ) • File event ( ) • Malware event ( ) Click the event icon next to the column set filtering options to filter the list of event fields according to the selected event type. **Note** Including many columns may degrade performance. You can view data for hidden columns by expanding an event row to view event details. • Reorder columns: Drag and drop the column heading. • Pin (freeze) columns to the left or right side of the table so they do not scroll: • Drag a column all the way to either left or right side of the table or drag and drop a column heading into the pinned area. • To unpin a column, drag the column out of the pinned area. • Resize columns. • Revert columns to the default setting. • Save column sets to quickly reload your customized view later. For more information, refer to Save a column set, on page 513 topic. Data is always sorted by time, with the most recent events on top.
Identify related events	Click a row to highlight other events that are related to this event. If needed, filter the events to display a small enough set of events. **Note** The initiator of a connection is not necessarily the same as the sender of a malware file. Search for the file or malware event associated with a connection event by filtering the unified events table with the **Source or Destination IP** filter.

To do this ...	Do this
View event details	Click the > (Expand) icon at the left end of the row. Event details do not include the field which has no data to display. **Tip** Alternatively, double-click on an event row to view the **Event Details** pane. When the **Event Details** pane is open, click on any event row in the table to load the details of that event.
Cross-launch to external resources	Click the ellipsis (⋮) in a table cell to see the options available for that cell value, if any. For more information, refer to Event investigation using web-based resources, on page 535.
Open multiple unified events windows	• Open unified events table in multiple browser tabs or windows to view different filtered data simultaneously. • Each new tab or window has the characteristics of the most recently modified tab or window. • To make any open tab or window as the template, make a minor change to it. • The system processes queries on multiple tabs sequentially. • Depending on the view (complex queries, or viewing in live view mode when the incoming event rate is high, for example), you may experience slower performance if more than 4 tabs are open simultaneously.
Save searches	Save custom searches as your favorites and quickly load them later. For more information, refer to Save a search in unified events, on page 512.
Bookmark or share query results	Bookmark or copy-paste the URL in the browser window. • The URL retrieves different events later if it uses the sliding time range. • The URL does not capture column visibility, size and order, and real-time streaming settings.

Set a time range in unified events

Set a time range in unified events to view firewall events for a specific period and control which events are displayed in the table.

When you change the time range, the unified events table automatically refreshes to reflect your changes. The time range that you select does not apply to other tables in the event viewer. For example, a time range that you select when viewing connection events does not apply to the unified events table and vice versa.

Note If your selected time range exceeds the event retention period allowed by your license, it will automatically adjust to fit within the retention period.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

By default, the unified events table displays events from the past hour.

Step 2 Click the current time range.

Step 3 Choose one of the following:

- If you want to see events for a fixed time range, click **Fixed Time Range** and choose the **Start time** and **End time**.

  To set the current time as the **End time**, click **Now**.

- If you want a sliding default time window (such as last one hour), select **Sliding Time Range** and specify the desired length.

  The table displays all the events generated from a specific start time—for example, the past hour—relative to the present. Refreshing the view ensures the window always displays events from the most recent hour of activity.

Step 4 Click **Apply**.

Filters in unified events

The **Unified Events** table displays firewall events from the past hour. Use these steps to filter and narrow the view for more granular analysis of your network traffic.

Filters help you quickly access critical information. For example, if you want to monitor application access for specific users, you can apply search criteria to isolate relevant firewall logs. The event viewer displays only the entries that match your criteria.

You can use both inclusion and exclusion criteria to refine your search results effectively.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Enter the filter criteria:

- To manually enter the filter criteria:

  a. Enter filter criteria in the search field, or select a filter from the drop-down list.

b. Enter the value for the selected filter criteria. Suggestions will appear in the drop-down list as you type.

- To pick the filter criteria from the table, click the dots in a cell and choose an option to include or exclude that value from your filter criteria.

  Tip
  - Use the **`Ctrl+click`** (Windows) or **`Command-click`** (Mac) key to quickly add an inclusion filter criteria.
  - Use the **`Alt+click`** (Windows) or **`Option-click`** (Mac) key to quickly add an exclusion filter criteria.

- Include operators (such as <, >, !) in the value field, preceding the value. For example, enter **`!Allow`** in the **Action** field to find all events with an action other than Allow.

Step 3 Perform the search.

Tip
You can use the **`Ctrl+Enter`** (Windows) or **`Command-Enter`** (Mac) key command to initiate a search.

Events in the unified events table are not aggregated when the displayed columns all hold identical values. Every event matching your filter criteria is listed individually.

The unified events table displays filtered results based on your criteria, showing only the events that match your inclusion and exclusion filters for more targeted analysis.

What to do next

To save a custom filter, see Save a search in unified events, on page 512 topic.

Save a search in unified events

Save custom searches as your favorites to quickly load them later for efficient firewall event analysis.

Note that this option is not available for the **Troubleshooting** table.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Click the **Events** tab.

Step 3 Enter a search criteria as described in the Filters in unified events, on page 511 topic.

Step 4 Click the **Favorite Searches** () icon on the search text box.

Step 5 Do one of the following:

- To save a new search, specify a search name and click **Save as new**.

- To overwrite a saved search, click **Edit** next to the saved search that you want to overwrite, and click **Overwrite**.

What to do next

To load a saved search, see Load a saved search in unified events, on page 513.

Load a saved search in unified events

If you have previously saved search criteria in **Unified Events**, you can quickly load the criteria and focus on particular firewall events without entering your criteria again.

Before you begin

Ensure you have already saved your preferred search criteria. For more information on saving search criteria, see Save a search in unified events, on page 512.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Click the **Favorite Searches** (☆) icon on the search text box.

Step 3 Click the saved search that you want to load.

Save a column set

Save custom column sets as your favorites to load them later or quickly toggle between custom tables.

This option allows you to create personalized table layouts for more efficient firewall event review. Note that this option is not available for the **Troubleshooting** table.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Click the column picker Icon (▥) and choose the set of columns that you want to save.

Step 3 Click the **Favorite column sets** (☆) icon.

Step 4 Do one of the following:

- To save a new column set, specify a column set name and click **Save as new**.

- To overwrite a favorite column set, click **Edit**(⋮) on the column set that you want to overwrite, and click **Overwrite**.

The custom column set is saved and can be loaded later for quick access to your preferred table layout.

What to do next

To load a saved column set, see Load a saved column set, on page 514 topic.

Load a saved column set

Apply preferred table layouts and streamline firewall event analysis by loading a previously saved column set in the **Unified Events** page.

Before you begin

Ensure you have already saved a column set. For more information on saving a column set, see Save a column set, on page 513.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Click the column picker icon (▥).

Step 3 Click the **Favorite column sets** (☆) .

Step 4 Click the column set that you want to load.

Unified events column details

Values in some field on the **Unified Events** page depend on the event type. See this table for values by event type for the default fields.

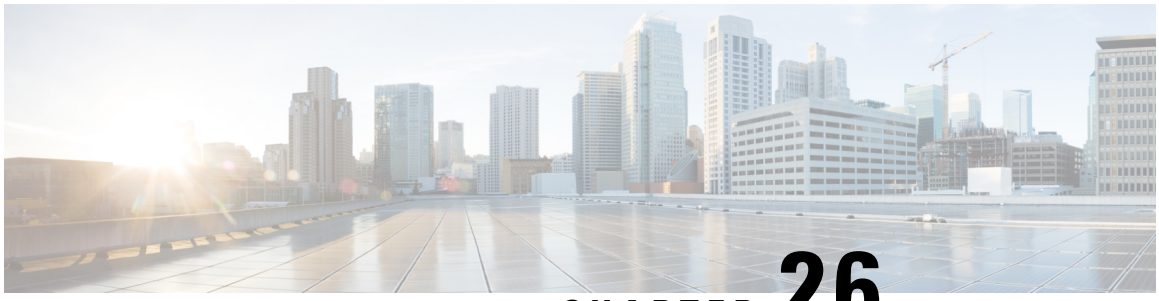
To see all event fields and their correspondences, use the column picker (▥) icon.

Unified events field	Connection or security-related connection event field	Intrusion event field	File event field	Malware event field
Time	First Packet	Time	Time	Time
Event Type	—	—	—	—
Action	Action	Inline Result	Action	Action

Unified events field	Connection or security-related connection event field	Intrusion event field	File event field	Malware event field
Reason	Reason	Reason	(Not applicable)	(Not applicable)
Source IP	Initiator IP	Source IP	Sending IP	Sending IP
Destination IP	Responder IP	Destination IP	Receiving IP	Receiving IP
Source Port/ICMP Type	Source Port	Source Port	Sending Port	Sending Port
Destination Port/ ICMP Type	Destination Port	Destination Port	Receiving Port	Receiving Port
Web Application	Web Application	Web Application	Web Application	Web Application
Rule	Access Control Rule	Access Control Rule	(Not applicable)	(Not applicable)
Policy	Access Control Policy	Intrusion Policy	File Policy	File Policy
Device	Device	Device	Device	Device



Note Even if logging is not enabled at the beginning of a connection, the system has and uses this value as the **Time** field in the unified events table. To check if a connection event was logged at the beginning and end of the connection, expand the event row for details. If both ends of the connection were logged, you will see a **Last Packet** field.



CHAPTER 26

Lookups

The following topics explain how to look up information about entities that may or may not be known to the system:

- Lookups in Cloud-Delivered Firewall Management Center, on page 517
- Lookup IP address details using Whois, on page 517

Lookups in Cloud-Delivered Firewall Management Center

In your Cloud-Delivered Firewall Management Center, use the lookup feature to find the Regional Information Registries (RIR) information (whois) for any IP address.

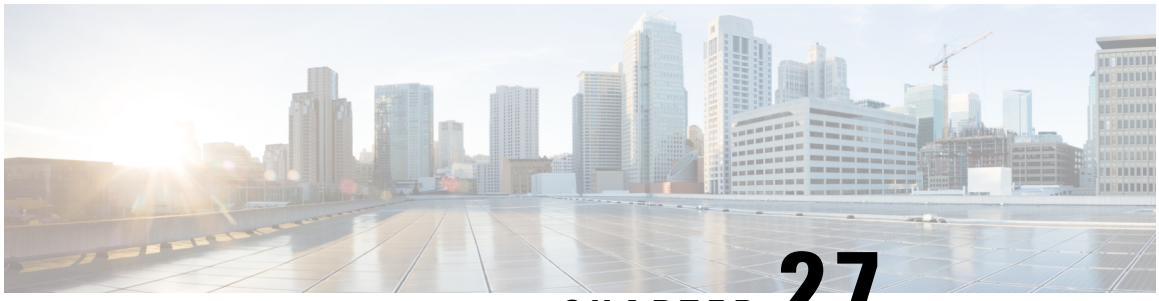
Lookup IP address details using Whois

Use Whois lookups to investigate the ownership and registration details of an IP address for security and network analysis.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Whois**.

Step 2 Enter an IP address you want to investigate and click **Search**.



CHAPTER 27

Event searches

An event search is a customizable query that

- retrieves specific information stored as events in database tables based on defined search criteria,
- can be saved with custom names and shared with other users or kept private for individual use, and
- returns only records that match the specified criteria for all fields in the search parameters.

Event search functionality

The system generates information that is stored as events in database tables. Events contain multiple fields that describe the activity that caused the appliance to generate the event. You can create and save searches customized for your environment for any of the different event types and save them to reuse later.

When you save a search you give it a name and specify whether the search will be available to you alone or to all users of the appliance. If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search. If you previously saved a search, you can load it, make any necessary modifications, and then start the search. Custom analysis dashboard widgets, report templates, and custom roles can also use saved searches. If you have saved searches, you can delete them from the Search page.

For some event types, the system provides predefined searches that serve as examples and can provide quick access to important information about your network. You can modify fields within the predefined searches for your network environment, then save the searches to reuse later.

The search criteria you can use depends on the type of search, but the mechanics are the same. Searches return only records that match the search criteria specified for all fields.

Note Searching a custom table requires a slightly different procedure.

- Search constraints, on page 520
- Perform a search, on page 524
- Save a search, on page 524
- Load a saved search, on page 524

Search constraints

A search constraint is a event data search feature that enables you to enter search criteria values to apply to fields defined for event tables and constrain the data displayed to your specific needs.

Each database table has its own search page where you can enter search constraint values to apply to fields defined for the table. Depending on the type of field, special syntax may be used to specify criteria such as wildcard characters or a range of numeric values.

Search results appear on workflow pages displaying each table field in columnar layout. Some database tables can additionally be searched using fields that are not displayed as columns on workflow pages. To determine whether such a constraint applies to your search results when viewing the results on a workflow page, click

Expand Arrow ( ) to view the active search constraints.

Best practice for general search constraints

When searching for events, consider these general guidelines:

- Many fields require wildcards for partial-match searches. All fields accept wildcards for these searches. For more information, refer to Wildcards and symbols in searches, on page 521.
- All fields accept negation (`!`).
- All fields accept comma-separated lists of search values. Records that contain any of the listed values in the specified field match that search criteria.
- All fields accept comma-separated lists enclosed in quotation marks as search values.
  - For fields that may contain only a single value, records with the specified field containing the exact string specified within the quotation marks match the search criteria. For instance, a search for `A, B, "C, D, E"` will match records where the specified field contains `"A"` or `"B"` or `"C, D, E"`. This permits matching on fields that include the comma in possible values.
  - For fields that may contain multiple values at the same time, records with the specified fields containing all of the values in the quote-enclosed comma-separated list match that search criteria.
  - For fields that may contain multiple values at the same time, search criteria may include single values as well as quote-enclosed comma-separated lists. For instance, a search for `A, B, "C, D, E"` on a field that may contain one of more of these letters matches records where the specified field contains `A` or `B`, or all of `C`, `D`, and `E`.
- Specify `n/a` in any field to identify events where information is not available for that field; use `!n/a` to identify the events where that field is populated.
- You can precede many numeric fields with greater than (`>`), greater than or equal to (`>=`), less than (**<**), less than or equal to (`<=`), equal to (`=`), or not equal to (`<>`) operators.

Note When you search a field with long complicated values (such as SHA-256 hash values), copy the search criteria value from source material and paste it into the appropriate field on the search page.

Wildcards and symbols in searches

When searching in all text fields in connection and Security Intelligence events and in most text fields in other event types, searches for partial matches in text fields require an asterisk (*) to represent unspecified characters in a string. Searches without an asterisk are exact-match searches in these fields. Even in fields that do not require wildcards, we recommend always using wildcards for partial-match searches.

For example, to find `example.com`, `www.example.com`, or `department.example.com`, search for `*.example.com`. Searching for `example.com` in most cases returns only `example.com`.

Search for non-alphanumeric characters

If you want to search for non-alphanumeric characters (including the asterisk character), enclose the search string in quotation marks. For example, to search for the string:

```
Find an asterisk (*)
```

enter:

```
"Find an asterisk (*)"
```

Objects and application filters in searches

The system allows you to create named objects, object groups, and application filters that can be used as part of your network configuration. You can use these objects, groups, and filters as search criteria when performing or saving searches.

When you perform a search, objects, object groups, and application filters appear in the format, `${object_name}`. For example, a network object with the object name `ten_ten_network` appears as `${ten_ten_network}` in a search.

Click the **Object** (+) icon that appears next to a search field to use an object as a search criterion.

Related Topics

Use the object manager, on page 2932

Time constraints in searches

The formats accepted by search criteria fields that take a time value are shown in the tables that follow.

Table 35: Time specification in search fields

Time formats	Example
`today [at HH:MMam|PM]`	`today` `today at 12:45pm`
`YYYY-DDMM-HH:MM:SS`	`2006-03-22 14:22:59`

You can precede a time value with one of the operators shown in this table:

Table 36: Time specification operators

Operator	Example	Explanation
<	`< 2006-03-22 14:22:59`	Returns events with a timestamp before 2:23 PM, March 22, 2006.
>	`> today at 2:45pm`	Returns events with a timestamp later than today at 2:45 PM.

IP addresses in searches

When specifying IP addresses in searches, you can enter an individual IP address, a comma-separated list of addresses, an address block, or a range of IP addresses separated with a hyphen (-). You can also use negation.

When you search for hosts by IP address, the results include all hosts for which at least one IP address matches your search conditions, that is, a search for an IPv6 address may return hosts whose primary address in IPv4.

When you use CIDR or prefix length notation to specify a block of IP addresses, the system uses only the portion of the network IP address specified by the mask or prefix length. For example, if you type `10.1.2.3/8`, the system uses `10.0.0.0/8`.

Table 37: Acceptable IP address syntax

To specify...	Type...	For example...
a single IP address	the IP address.	`192.168.1.1` `2001:db8::abcd`
multiple IP addresses using a list	a comma-separated list of IP addresses. Do not add a space before or after the commas.	`192.168.1.1,192.168.1.2` `2001:db8::b3ff,2001:db8::0202`
a range of IP addresses that can be specified with a CIDR block or prefix length	the IP address block in IPv4 CIDR or IPv6 prefix length notation.	`192.168.1.0/24` This specifies any IP in the 192.168.1.0 network with a subnet mask of 255.255.255.0, that is, 192.168.1.0 through 192.168.1.255.
a range of IP addresses that cannot be specified with a CIDR block or prefix	the IP address range using a hyphen. Do not add a space before or after the hyphen.	`192.168.1.1-192.168.1.5` `2001:db8::0202-2001:db8::8329`

negation of any of the other ways to specify IP addresses or ranges of IP addresses	an exclamation point in front of the IP address, block, or range.	`192.168.0.0/32,!192.168.1.10`
		`!2001:db8::/32`
		`!192.168.1.10,!2001:db8::/32`

Related Topics

IP address conventions

URLs in searches

When searching for URLs, include wildcards. For example, use **`*example.com*`** to find all variations of the domain, such as **`https://example.com`** and **`division.example.com`** and **`example.com/division/`**.

Managed devices in searches

When you group devices—whether just on the Cloud-Delivered Firewall Management Center, or as actual high availability or scalability configurations—searching for the name for the group correctly returns results for all devices in the group.

If the system finds a match for a group, it replaces the group name with the appropriate member device names for the purpose of performing the search. When you save a search that uses a device group in the device field the system saves the name specified in the device field and performs the device name replacement again each time the search is executed.

Ports in searches

The Cloud-Delivered Firewall Management Center accepts specific syntax for port numbers in searches.

You can enter port numbers in these formats:

- A single port number.
- A comma-separated list of port numbers.
- Two port numbers separated by a dash to represent a range of port numbers.
-
- A port number or range of port numbers preceded by an exclamation mark to indicate a negation of the specified ports.

Note Avoid using spaces when specifying port numbers or ranges.

Table 38: Port syntax examples

Example	Description
`21`	Returns all events on port 21, including TCP and UDP events.
`!23`	Returns all events except those on port 23.
`25/tcp`	Returns all TCP-related intrusion events on port 25.
`21/tcp,25/tcp`	Returns all TCP-related intrusion events on ports 21 and 25.
`21-25`	Returns all events on ports 21 through 25.

Perform a search

Search for specific events or data records within the system to analyze health events, and other types of monitored data.

Procedure

Step 1 Click **Search** from a workflow page.

Step 2 From the table drop-down list, select the type of event or data to search.

Step 3 Enter your search criteria in the appropriate fields.

Step 4 If you want to use the search again in the future, save the search as described in Save a search, on page 524.

Step 5 Click **Search** to start the search.

Your search results appear in the default workflow for the table you are searching, constrained by time (if applicable).

Save a search

Save search criteria for future use to streamline repeated searches and maintain consistent filtering parameters.

Before you begin

- Establish search criteria as described in Perform a search, on page 524, or load a saved search as described in Load a saved search, on page 524.

Procedure

Step 1 From the Search page, if you want to save the search as private so only you can access it, check the **Private** checkbox.

Step 2 You have these options:

- To save a new version of a loaded search, click **Save As New**.
- To save a new search, or overwrite a custom search using the same name, click **Save**.

Load a saved search

Load a saved search to quickly access previously configured search criteria without having to recreate the search constraints.

Procedure

Step 1 Click **Search** from a workflow page.

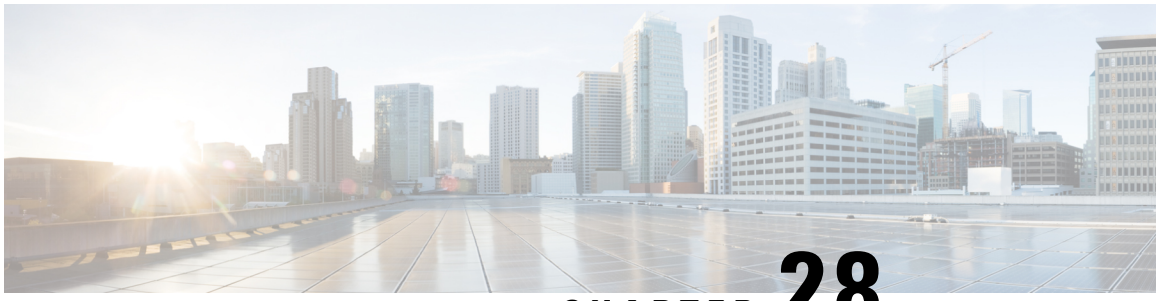
Step 2 From the table drop-down list, choose the type of event or data to search.

Step 3 Choose the search you want to load from the **Custom Searches** list or the **Predefined Searches** list.

Step 4 If you want to use different search criteria, change the search constraints.

Step 5 If you want to use a changed search again in the future, save the search as described in Save a search, on page 524.

Step 6 Click **Search**.



CHAPTER 28

Send Firewall Threat Defense Events to Cisco Splunk

- Splunk integration: Send events directly from Cloud-Delivered Firewall Management Center , on page 527
- Guidelines and limitations for Splunk integration, on page 527
- Configure Splunk in Cloud-Delivered Firewall Management Center , on page 528
- Configure Secure Firewall App in Splunk, on page 533

Splunk integration: Send events directly from Cloud-Delivered Firewall Management Center

Cisco Splunk is a Security Information and Event Management (SIEM) tool that provides visibility and monitoring of security events across Cisco Secure Firewall devices.

In Cloud-Delivered Firewall Management Center versions earlier to 10.0, security events were sent to Splunk using eStreamer. From Cloud-Delivered Firewall Management Center version 10.0, you can send events to Splunk from the Firewall Threat Defense device. The wizard-driven interface helps you set up Splunk integration effortlessly.

This integration allows you to perform these actions.

- Customize event flow by specifying event types—such as connection, intrusion, malware, file, user activity, correlation, discovery, intrusion packet—and their sources (Firewall Threat Defense device) according to your monitoring requirements.
- Select the source interface for sending syslog events. You can choose to send events from the Firewall Threat Defense device's management or data interfaces.
- Create profiles with various configurations to suit different monitoring requirements.

Guidelines and limitations for Splunk integration

This topic lists the guidelines and limitations that apply to Splunk integration for event forwarding and syslog processing.

- When you enable Splunk integration, your existing syslog server configuration is not migrated.
- When migrating to Cloud-Delivered Firewall Management Center, you must repeat the Splunk integration procedure because the Splunk configuration from the On-Prem Firewall Management Center does not get migrated.
- If you set up domains in Cloud-Delivered Firewall Management Center, you can configure Splunk only from leaf domains.
- You can create a maximum of 15 Splunk profiles per domain. However, multiple domains can be configured to send syslog messages to the same or different servers.
- Do not create duplicate profiles for the same devices because it results in devices sending duplicated events to the Splunk server.
- Events are duplicated in Splunk destinations, under any of the following scenarios:
  - Platform Settings are configured with at least one syslog configuration.
  - Splunk profile is enabled using at least one data interface option for events.
- Only basic TLS encryption is supported over data interfaces from devices; client and server authentication are not supported.
- The syslog payload is sent to the Splunk server in JSON format.
- In a high availability set up, Splunk configuration is synchronized between the Cloud-Delivered Firewall Management Center HA pair. Only active unit sends the events. In the case of a failover, the new active unit will start sending the events.
- The message header of syslog events sent through data interfaces contains only the syslog tag and not the device IP address.
- cdFMC does not support sending events from a Cloud-Delivered Firewall Management Center source. Hence, cdFMC cannot have events that support only Cloud-Delivered Firewall Management Center source.

Configure Splunk in Cloud-Delivered Firewall Management Center

The Splunk integration wizard in Secure Firewall Management Center allows you to create profiles that send events and syslog data to a designated server. You can create multiple profiles to configure any number of servers for different combinations of devices and event types.

Before you begin

- Ensure that the Splunk server can be reached by Firewall Threat Defense device.
- The Cisco Secure Firewall App in Splunk does not support TLS. Hence, if you choose to use the TLS protocol to send events to Splunk, configure TLS on the Splunk server. For TLS configuration instructions, refer to the section *Configure Splunk indexing and forwarding to use TLS certificates* under *Manage Users and Security* in the Splunk Administer guide.

- Create required objects such as host, security zone, interface group, certificate, and so on, before starting the configuration procedure. Although you can navigate from the **Splunk integration** wizard to create objects, having them in advance will provide a smoother integration experience.

The Splunk integration wizard allows you to create a profile that enables you to stream events and syslog from managed devices to a specific server.

You can create multiple profiles to configure any number of servers for various combinations of devices and events. For example, you can create multiple profiles to send a specific set of events to one server and all the remaining events to a different server. Each profile is independent, but they all apply additively.

Procedure

	Command or Action	Purpose
Step 1	To open the **Splunk integration** wizard, go to **Integration** > **Splunk**.	
Step 2	The steps for configuring Splunk integration in Cloud-Delivered Firewall Management Center are listed in this table.	

	Do This	More Information
Step 1	Configure Splunk or similar SIEM tool server.	Refer to Configure Splunk server, on page 530.
Step 2	Choose the event types that you want to send to the Splunk server.	Refer to Select event types, on page 530.
Step 3	Specify the devices and the interfaces from which you want to send syslog events to Splunk.	Refer to Select devices and interfaces, on page 531.
Step 4	(Optional) Specify the device certificate to be used for sending events securely to Splunk.	Refer to Configure Firewall certificates for TLS, on page 532.
Step 5	View the summary of the profile that is being created.	Refer to Review and submit the Splunk profile, on page 532.

Configure Splunk server

Set up a Splunk syslog server to efficiently collect and manage log events from Secure Firewall devices. Use this task when adding a Splunk server as an external syslog destination for your Secure Firewall deployment.

Procedure

Step 1 In the **Configure Splunk server** page, from the **Host object or IP address** drop-down, choose the host object or enter an IP address of the Splunk server. To create a host object, click the **Create** link in the drop-down list.

Step 2 Click the **Protocol** (UDP, TCP, or TLS) that you want the Splunk server to use for communicating with Cloud-Delivered Firewall Management Center and Firewall Threat Defense device.

Step 3 In the **Port** field, enter the port number applicable for the selected protocol:

- UDP: Enter the port number 514, or a number between 1025 and 65535. The default UDP port is 514.
- TCP: Enter the port number between 1025 and 65535. The default TCP port is 1470.
- TLS: Enter the port number between 1 and 65535. The default TLS port is 6514.

Step 4 (Optional) If you selected **TLS**, specify the trusted CA name to securely establish the connection on the Splunk server. From the **Trusted certificate authority** drop-down list, choose the CA. (To import a new CA object, click the **Create** link in the drop-down list.)

Step 5 To forward the events logged by a facility, from the **Facility** drop-down list, choose the corresponding facility.

Step 6 To forward the events with a specific severity, from the **Severity** drop-down list, choose the severity level.

Step 7 (Optional) In the **Tag** field, enter an alphanumeric string by which to identify the message in the server.

Step 8 (Optional) To cancel the creation of the Splunk profile, click **Cancel**.

Step 9 To move to the next step in the Splunk configuration, click **Next**.

Select event types

You can specify the type of events that you want to send to Splunk. The recommended event types are selected by default. In Cloud-Delivered Firewall Management Center , the source for event types is always **FTD**.

Procedure

Step 1 Select the check box to specify the event types.

Step 2 After selection of the event types, to reset to default settings, click **Revert to system defaults**.

Step 3 (Optional) To cancel the creation of the Splunk profile, click **Cancel**.

Step 4 (Optional) To go back to the previous step, click **Back**.

Step 5 To move to the next step in the Splunk configuration, click **Next**.

Select devices and interfaces

You can specify the Firewall Threat Defense devices and interfaces from which you want to send syslog events to Splunk. To send events from multiple interfaces, use security zones or interface groups.

Note If you are using security zones or interface groups, only routed, management, switched, and loopback zones and groups are allowed.

Procedure

Step 1 Under **Select devices**, click the relevant options:

- **Use management interface** : Click this button to configure the management interface of Firewall Threat Defense device to send the events to Splunk. When this option is chosen, all the devices in the current domain also receive the Splunk configuration.
- **Use security zones and interface groups to specify devices and interfaces** : Click this button to configure the interfaces of corresponding devices in the selected security zones and interface groups to send their respective events to Splunk.
  - From the **Security zones and interface groups** drop-down list, choose the required zones and groups.
  - To create a security zone or interface group for the device's interfaces, click **Create** in the **Certificates** drop-down list.
- **Manually select devices and interfaces** : Click this button to send the events from the deployed device to Splunk. From the **Interface** drop-down list, choose the configured interface through which you want the events to be sent to Splunk. You can choose only security zones that include this device.

  Note
  - Virtual Tunnel Interface (VTI) and Dynamic Virtual Tunnel Interface (DVTI) are excluded from Splunk configuration.
  - To encrypt syslog events over TLS, use the management interface to add a certificate. Encryption is not supported for data interfaces configured for TLS.
  - Click the **Send events from this device** toggle button corresponding to the device and its selected interface.
  - To create a security zone or interface group for the device's interfaces, click **Create**.

Step 2 (Optional) To cancel the creation of the Splunk profile, click **Cancel**.

Step 3 (Optional) To go back to the previous step, click **Back**.

Step 4 To move to the next step in the Splunk configuration, click **Next**.

Configure Firewall certificates for TLS

You can specify the device certificate to be used for sending events securely to Splunk. This step is optional. It is applicable only if you use the TLS protocol to send events to Splunk.

Before you begin

You can configure the certificates when you have set TLS protocol while configuring the Splunk server. Add certificate object for these clients:

- For Cloud-Delivered Firewall Management Center, use internal certificate objects for authentication.
- For Firewall Threat Defense, use a certificate enrollment object. First, create the object at **Objects** > **Object Management** > **PKI** > **Certificate Enrollment**. Then, enroll the certificate on a device at **Devices** > **Certificates**.

Procedure

Step 1 From the **Firewall Management Center client certificate** drop-down list, choose the certificate. To create an internal certificate, click **Create**.

Step 2 In the **Firewall Threat Defense certificates** box, use the search box to choose the device whose certificate you want to assign.

Step 3 From the **Certificates** drop-down list, choose the certificate. Only successfully enrolled device certificates are listed. To enroll a new device certificate, click **Create**.

Step 4 (Optional) To cancel the creation of the Splunk profile, click **Cancel**.

Step 5 (Optional) To go back to the previous step, click **Back**.

Step 6 To move to the next step in the Splunk configuration, click **Next**.

Review and submit the Splunk profile

You can view the selections made for the Splunk configuration profile.

Procedure

Step 1 Under **Profile Name**, a system-generated name is displayed. Edit the value in the **Name** field, if you want to use a different name.

Note
The name must start with a letter, a number, or an underscore (_). The name can contain include letters, numbers, and special characters: period (.), hyphen (-), underscore (_), and plus (+).

Step 2 To change a value from a previous Splunk server configuration, click **Edit** next to the attribute.

Step 3 (Optional) To cancel the creation of the Splunk profile, click **Cancel**.

Step 4 (Optional) To go back to the previous step, click **Back**.

Step 5 To save the Splunk profile, click **Submit**.

Step 6 If you have configured Threat Defense events to be sent to Splunk, click **Deploy**.

Configure Secure Firewall App in Splunk

To ensure that the Splunk server is reachable to receive the events, configure the Cisco Secure Firewall App in Splunk.

Before you begin

Ensure that you have obtained a Splunk server license and a Cisco Security Cloud account.

Procedure

Step 1 Download and install the Splunk server using the instructions provided in Splunk Enterprise Installation Manual.

Step 2 To install Splunk license, log in to your Splunk server's web interface.

Step 3 Go to **Settings** > **Licensing** > **Add license**. Use the license received by email.

Note
Make sure to restart Splunk to complete license registration.

Step 4 Download Cisco Security Cloud from Splunkbase Apps.

Step 5 To install Cisco Security Cloud, log in to your Splunk server's web interface.

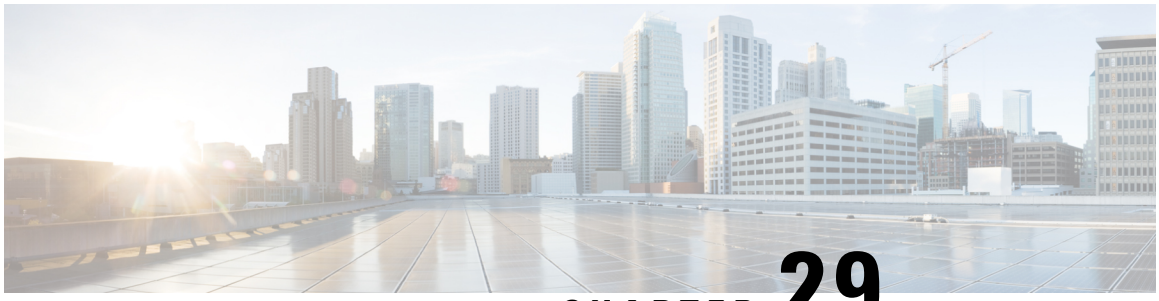
Step 6 Go to **Apps** > **Manage Apps** > **Install app from file**.

Step 7 Click **Browse** and select the downloaded application file and upload.

Step 8 To configure the server to receive the syslog events from the Firewall Threat Defense device, go to **Apps** > **Cisco Security Cloud** > **Secure Firewall** > **Configure Application**, and then click the **Syslog** tab.

- In the **Input Name** field, enter any name.
- In the **Input Type** field, enter UDP or TCP.

  TLS is not supported in the Cisco Security Cloud application. Use rsyslog or syslog-ng to establish TLS on Splunk server. For detailed procedures, refer to Configure TLS on Splunk Server.
- In the **Port** field, enter a value between 1025 and 65535. The default port is 514.
- Leave the **Host** field blank. This will allow the Splunk to process events from any Firewall Threat Defense.
- In the **Source Type** field, enter *cisco:ftd:syslog*.
- In the **Interval** field, enter *600* seconds.



CHAPTER 29

Event Investigation Using Web-Based Resources

• Event investigation using web-based resources, on page 535

Event investigation using web-based resources

Event investigation using web-based resources is a security analysis method that

- enables quick access to threat intelligence from external sources outside of the Cloud-Delivered Firewall Management Center,
- allows direct navigation from events to relevant information in external resources based on IP addresses, ports, protocol, domain, and SHA-256 hash, and
- provides access to predefined links to commonly used Cisco and third-party threat intelligence services, plus custom links to other web-based services and SIEMs.

Event investigation capabilities

The contextual cross-launch feature supports various investigation scenarios:

- Look up a suspicious source IP address in a Cisco or third-party cloud-hosted service that publishes information about known and suspected threats
- Look for past instances of a particular threat in your organization's historical logs, if your organization stores that data in a Security Information and Event Management (SIEM) application
- Look for information about a particular file, including file trajectory information, if your organization has deployed Cisco Secure Endpoint

When investigating an event, you can click directly from an event in the event viewer or dashboard in the Cloud-Delivered Firewall Management Center to the relevant information in the external resource. This lets you quickly gather context around a specific event based on its IP addresses, ports, protocol, domain, and/or SHA-256 hash.

Note Some resources may require an account or a product purchase.

Talos IP address investigation

Suppose you are looking at the Top Attackers dashboard widget and you want to find out more information about one of the source IP addresses listed. You want to see what information Talos publishes about this IP address, so you choose the "Talos IP" resource. The Talos website opens to a page with information about this specific IP address.

Manage contextual cross-launch resources

Manage external web-based resources using the **Analysis** > **Advanced** > **Contextual Cross-Launch** page.

Pre-defined resources offered by Cisco are marked with the Cisco logo. The remaining links are third-party resources.

You can disable or delete any resources that you do not need, or you can rename them, for example by prefixing a name with a lower-case "z" so the resource sorts to the bottom of the list. Disabling a cross-launch resource disables it for all users. You cannot reinstate deleted resources, but you can re-create them.

To add a resource, refer to Add contextual cross-launch resources, on page 536.

Requirements for custom contextual cross-launch resources

A custom contextual cross-launch resource is an integration feature that

- enables web browser access to predefined or custom resources, and
- restricts configuration and usage based on user role permissions.

Custom contextual cross-launch resource specifications

Custom contextual cross-launch resources must meet these requirements:

- Resources must be accessible via web browser.
- Only HTTP and HTTPS protocols are supported.
- Only GET requests are supported; POST requests are not.
- Encoding of variables in URLs is not supported. While IPv6 addresses may require colon separators to be encoded, most services do not require this encoding.
- Up to 100 resources can be configured, including pre-defined resources.

Add contextual cross-launch resources

In multidomain deployments, you can see and use resources in parent domains, but you can only create and edit resources in the current domain. The total number of resources across all domains is limited to 100.

Before you begin

- If you are adding links to a Secure Network Analytics appliance, check to see if the links you want already exist; most links are automatically created for you when you configure Security Analytics and Logging (On Premises).

- Refer to Requirements for custom contextual cross-launch resources, on page 536.
- If needed for the resource you will link to, create or obtain an account and the credentials needed for access. Optionally, assign and distribute credentials for each user who needs access.
- Determine the syntax of the query link for the resource that you will link to:

  Access the resource via browser and, using the documentation for that resource as needed, formulate the query link needed to search for a specific sample of the type of information you want your query link to find, such as an IP address.

  Run the query, then copy the resulting URL from the browser's location bar.

  For example, you might have the query URL **`https://www.talosintelligence.com/reputation_center/lookup?search=10.10.10.10`**.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Contextual Cross-Launch**.

Step 2 Click **New Cross-launch**.

In the form that appears, all fields marked with an asterisk require a value.

Step 3 Enter a unique resource name.

Step 4 Paste the working URL string from your resource into the **URL Template** field.

Step 5 Replace the specific data (such as an IP address) in the query string with an appropriate variable: Position your cursor, then click a variable (for example, **IP**) once to insert the variable.

In the example from the "Before You Begin" section above, the resulting URL might be **`https://www.talosintelligence.com/reputation_center/lookup?search={IP}`**. When the contextual cross-launch link is used, the {IP} variable in the URL will be replaced by the IP address that the user right-clicks on in the event viewer or dashboard.

For a description of each variable, hover over the variable.

You can create multiple contextual cross-launch links for a single tool or service, using different variables for each.

Step 6 Click **Test with example data** ( ) to test your link with example data.

Step 7 Fix any problems.

Step 8 Click **Save**.

Investigate events using contextual Cross-Launch

Contextual cross-launch allows you to right-click on event data and immediately access external analysis tools and threat intelligence resources. The available options depend on the data type you select, ensuring you only see relevant investigation resources.

Before you begin

If the resource you will access requires credentials, make sure you have those credentials.

Follow these steps to investigate events using contextual cross-launch:

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, click **Analysis** > **Unified Events**.

Step 2 Right-click the event of interest and choose the contextual cross-launch resource to use.

If necessary, scroll down in the context menu to see all available options.

The data type you right-click on determines the options you see; for example, if you right-click an IP address, you will only see contextual cross-launch options that are relevant to IP addresses.

For example, to get threat intelligence from Cisco Talos about a source IP address in the intrusion event, choose **Talos SrcIP** or **Talos IP**.

If a resource includes multiple variables, the option to choose that resource is available only for events that have a single possible value for each included variable.

The contextual cross-launch resource opens in a separate browser window.

It may take some time for the query to be processed, depending on the amount of data to be queried, speed of and demand on the resource, and so on.

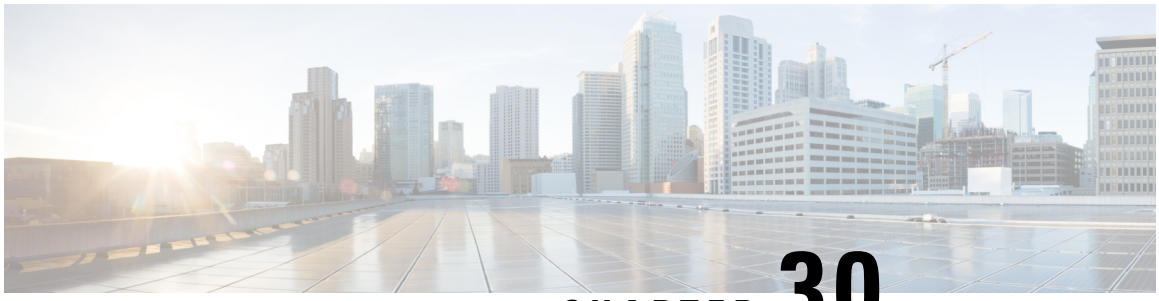
Step 3 Sign in to the resource if necessary.



PART X

Events and Assets

- Monitor Cloud-Delivered Firewall Management Center-Managed Threat Defense Device Events, on page 541
- FTD Dashboard, on page 559
- Connection Logging, on page 565
- Advanced Logging Syslog Fields, on page 587
- User Identity and Activity, on page 595



CHAPTER 30

Monitor Cloud-Delivered Firewall Management Center-Managed Threat Defense Device Events

View events from Threat Defense devices that are managed by Cloud-Delivered Firewall Management Center in Security Cloud Control.

- About Security Analytics and Logging, on page 541
- Comparison of SAL Remote Event Storage and Monitoring Options, on page 542
- About SAL (OnPrem), on page 542
- Manage SAL (OnPrem) for Security Cloud Control-Managed Firewall Threat Defense Devices, on page 543
- Configure SAL (OnPrem) Integration, on page 545
- About Security Analytics and Logging(SaaS) for Firewall Threat Defense, on page 548
- Configure the SAL (SaaS) Integration, on page 548

About Security Analytics and Logging

Security Analytics and Logging (SAL) is a central log management and advanced threat detection service which delivers scalable Cisco firewall logging and correlated analytics. Central logging helps in providing visibility, helps troubleshoot network access issues including disruptions, and enables device and overall network health monitoring. Analytics provide detection against advanced threats.

The SAL service is available in the following two methods:

- Security Analytics and Logging (SaaS)—A hosted software as a service (SaaS) which stores events and provides data for security analytics. This service connects the Security Analytics and Logging cloud data store to the firewall cloud manager, Security Cloud Control (Security Cloud Control).

  In this documentation, this method is also referred to as SAL (SaaS).

- Security Analytics and Logging (On Premises)—A service that runs on the Secure Network Analytics appliances to store event logs at the customer's own premises. This service connects the Security Analytics and Logging (On Premises) data to the on-premises manager, Secure Firewall Management Center.

  In this documentation, this method is also referred to as SAL (OnPrem).

For more information about Security Analytics and Logging, see https://www.cisco.com/c/en/us/products/security/security-analytics-logging/index.html.

Comparison of SAL Remote Event Storage and Monitoring Options

SAL integration shows similar options for storing event data externally to a Firewall Management Center and Security Cloud Control:

	SAL (OnPrem)	SAL (SaaS)
Why choose this solution?	You want to increase your on-premises firewall event data storage capacity, retain this data for a longer period, and export your event data to the Secure Network Analytics appliance.	You want to send firewall events to the Cisco cloud for storage and optionally make your firewall event data available for security analytics using integrated cloud services such as Cisco XDR.
Licensing	Purchase license and set-up the storage system behind your firewall. For more information, see Licensing for SAL (OnPrem), on page 543	Purchase license and a data storage plan and send your data to the Cisco cloud. For more information, see Licensing for SAL (SaaS), on page 548
Supported event types	• Connection • File and Malware • Intrusion • LINA • Security Intelligence	• Connection • File and Malware • Intrusion • Security Intelligence
Supported methods to send events	Supports both, syslog and direct integration.	Supports both, syslog and direct integration.
Event viewing	• View events on the Secure Network Analytics Manager. • Cross-launch from the Firewall Management Center event viewer to view events on the Secure Network Analytics Manager. • View remotely stored connection and security intelligence events in the management center.	View events in Security Cloud Control or Secure Network Analytics Manager, depending on your license. Cross-launch from the Firewall Management Center event viewer.

About SAL (OnPrem)

You can configure SAL (OnPrem) to store firewall event data for increased storage at a larger retention period. By deploying Secure Network Analytics appliances and integrating them with your firewall deployment, you can export your event data to a Secure Network Analytics appliance.

This provides you with the following capabilities:

- Store events on the Secure Network Analytics appliance.
- Specify this remote data source to view these events in the management center.
- Review event data from the Secure Network Analytics Manager (formerly Stealthwatch Management Console) Web App UI using the *Event Viewer*.
- Cross-launch from the management center UI to the *Event Viewer* to view additional context on the information from which you cross-launched.

Licensing for SAL (OnPrem)

You must obtain the Logging and Troubleshooting smart license to use SAL (OnPrem). You can obtain the license based on the amount of data you anticipate while sending the syslog data from your firewall deployment to your Secure Network Analytics appliance on a daily basis.

For information on licensing the Secure Network Analytics appliances, see Secure Network Analytics Smart Software Licensing Guide.

For information on the available SAL (OnPrem) licensing options, see the Cisco Security Analytics and Logging Ordering Guide.



Note For license calculation purposes, the amount of data is reported to the nearest whole GB. For example, If you send 4.9 GB in a day, it is reported as 4 GB.

Manage SAL (OnPrem) for Security Cloud Control-Managed Firewall Threat Defense Devices

Starting with Secure Firewall Threat Defense(formerly Firepower Threat Defense) version 7.2, you can choose to send fully qualified events that are generated by Security Cloud Control-managed Firewall Threat Defense devices to the Firewall Management Center. The Firewall Management Center receives and displays data analytics for these events. The Firewall Management Center receiving and displaying the event data is also referred to as an analytics-only management center. .

If your devices are enabled to send connection events to a Secure Network Analytics Manager using SAL (OnPrem), you can view and work with these remotely stored events in the management center event viewer and context explorer, and include them when generating reports. By deploying the Secure Network Analytics appliance and integrating it with the firewall deployment, you can export the event data to the Secure Network Analytics appliance. This allows you to view and manage the events in the management center UI. From the management center interface, you can also cross-launch to Secure Network Analytics Manager to view and manage the events data.

The management center can receive and display event analytics for the following Security Cloud Control-managed Firewall Threat Defense devices:

- New or existing Firewall Threat Defense devices onboarded to Security Cloud Control

For information on onboarding a Firewall Threat Defense device to Security Cloud Control, see Prerequisites to Onboard a Device to Cloud-Delivered Firewall Management Center, on page 21.

The workflow is as follows:

1. Onboard a Firewall Threat Defense device to Security Cloud Control.

   Onboard the Firewall Threat Defense devices using the onboarding methods that are described in Prerequisites to Onboard a Device to Cloud-Delivered Firewall Management Center, on page 21. The onboarding process includes assigning policies and choosing the appropriate licenses.

2. Register this Firewall Threat Defense device in the appropriate management center.

   For the management center to display events generated by a Security Cloud Control-managed Firewall Threat Defense device, you must register the Firewall Threat Defense device in the management center. To register this device in the Firewall Management Center, enable the device to be registered using the **configure manager add** {*hostname* | *IPv4_address* | *IPv6_address*}*reg_key*[*nat_id*] CLI, and then add the device to the Firewall Management Center using the **Security Cloud Control Managed Device** check box.



Note The registration key and the NAT ID must be unique from those used while onboarding the device to Security Cloud Control.

   For more information, see *Add a Device to the Management Center* and *Complete the Threat Defense Initial Configuration Using the CLI* in Firepower Management Center Device Configuration Guide.

3. View events in the management center or cross-launch to a configured Secure Network Analytics Manager.

   To view and work with the events in the management center event viewer. If the Secure Network Analytics appliance is deployed and integrated with the firewall deployment, you can export the event data to the Secure Network Analytics appliance. This allows you to cross-launch from the management center UI to the Secure Network Analytics Manager to view and manage the events data.

   For more information, see *Events and Assets* and *Event Analysis Using External Tools*.

- Existing Firewall Threat Defense devices on the management center.

  You can change the management of the Firewall Threat Defense devices from management center to Security Cloud Control using the change threat defense manager functionality. The change threat defense manager functionality provides you to ability to change the management of Firewall Threat Defense devices from management center to Security Cloud Control. While changing the manager, you can choose to retain the events data generated by these threat defense devices on the management center. If you choose to retain the events data on the management center, a copy of the Firewall Threat Defense device in an analytics-only mode is retained on the management center.

  For more information, see Migrate Secure Firewall Threat Defense to Cloud.

  The workflow is as follows:

  1. Onboard the management center to Security Cloud Control

     To onboard the existing Firewall Threat Defense devices from management center to Security Cloud Control, you must onboard the appropriate management center to Security Cloud Control.

For more information, see Onboard an FMC.

2. Complete the change threat defense management process

   During the change threat defense management process, while changing the device manager, you can choose to retain events data generated by these Firewall Threat Defense devices on the management center.

   For more information, see Migrate Secure Firewall Threat Defense to Cloud.

3. View events in the management center or cross-launch to configured Secure Network Analytics appliance.

   To view and work with the events in the management center event viewer. If the Secure Network Analytics appliance is deployed and integrated with the firewall deployment, you can export the event data to the Secure Network Analytics appliance. This allows you to cross-launch from the management center UI to the Secure Network Analytics Manager to view and manage the events data.

   For more information, see Events and Assets and Event Analysis Using External Tools.

Configure SAL (OnPrem) Integration

Configure Security Cloud Control to send events to the Secure Network Analytics appliance.

Deploy a Cisco Secure Network Analytics Flow Collector to receive events, a Cisco Secure Network Analytics Data Store (containing 3 Cisco Secure Network Analytics Data Nodes) to store events, and a manager. The Firewall Threat Defense devices send event data to the flow collector from where the events are sent to the Data Store for storage. From the Cloud-Delivered Firewall Management Center user interface, you can cross-launch the manager to view more information about the stores events.

Starting with Firewall Threat Defense Version 7.2, you can choose to associate different flow collectors to different devices.

Configure a Secure Network Analytics Manager

Configure the Secure Network Analytics Manager deployment to integrate SAL (OnPrem) with Security Cloud Control-managed Firewall Threat Defense devices.

Before you begin

Ensure the following:

- You have a provisioned Security Cloud Control tenant and have the following Security Cloud Control user roles:
  - Admin
  - Super admin
- Your Firewall Threat Defense devices are working as expected and are generating events.
- If you are currently using syslog to send events to the Secure Network Analytics Manager from device versions that support sending events directly, disable syslog for those devices (or assign those devices

an access control policy that does not include syslog configurations) to avoid duplicating events on the remote volume.

- You have the hostname or the IP address of your Secure Network Analytics Manager.



Note You may be logged out of the Secure Network Analytics Manager during the registration process; complete any work in progress before you start with the deployment wizard.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 From the Security Cloud Control menu, navigate **Administration** > **Firewall Management Center** to open the **Services** page.

Step 3 Select **Cloud-Delivered FMC** and then click **Configuration**.

Step 4 Navigate to **Integration** > **Security Analytics & Logging**.

Step 5 In the **Secure Network Analytics Manager Only** widget, click **Start**.

Step 6 Enter the hostname or the IP address and port number of the Secure Network Analytics Manager and click **Next**.

Step 7 Deploy the changes to the managed devices.

The event data is not logged to the SAL (OnPrem) until the logging policy changes are deployed to the registered Firewall Threat Defense devices.

Note
If you must change any of these configurations, run the wizard again. If you disable the configuration or run he wizard again, all settings except the account credentials are retained.

You can view and work with these remotely stored events in the event viewer and context explorer in the management center, and include them when generating reports. You can also cross-launch from an event in the management center to view related data on your Secure Network Analyticss appliance.

For more information, see the online help for the management center.

Step 8 Click **OK**.

Configure a Secure Network Analytics Data Store

Configure a Secure Network Analytics data store deployment to integrate SAL (OnPrem) with Firewall Threat Defense devices that are Security Cloud Control-managed.

Before you begin

Ensure the following:

- You have a provisioned Security Cloud Control tenant and have the following Security Cloud Control user roles:

- Admin
- Super admin

- Your Firewall Threat Defense devices are working as expected and generating events.
- If you are currently using syslog to send events to the Secure Network Analytics appliance from device versions that support sending events directly, disable syslog for those devices (or assign those devices an access control policy that does not include syslog configurations) to avoid duplicate events on the remote volume.
- Gather the following information:
  - The hostname or the IP address of your Secure Network Analytics Manager.
  - The IP address of your flow collector.

Note You may be logged out of the Secure Network Analytics Manager during the registration process; complete any work in progress before you start with the deployment wizard.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 From the Security Cloud Control menu, navigate **Administration** > **Firewall Management Center** to open the **Services** page.

Step 3 Choose **Cloud-Delivered FMC** and click **Configuration**.

Step 4 Navigate to **Integration** > **Security Analytics & Logging**.

Step 5 In the **Secure Network Analytics Data Store** widget, click **Start**.

Step 6 Enter the hostname or the IP address and port number of the flow collector.

To add more flow collectors, click **+Add another Flow Collector**.

Step 7 If you have configured more than one flow collector, associate the managed devices with different flow collectors:

Note
By default, all the managed devices are assigned to the default flow collector.

a) Click **Assign Devices**.
b) Select the managed devices that you want to assign.
c) From the reassign device drop-down list, choose the flow collector.

If you do not want a managed device to send event data to any of the flow collectors, select that device, and choose **Do not log to flow collector from the reassign device drop-down list.**

You can change the default flow collector by hovering over the intended flow collector and clicking **Set default**.

d) Click **Apply Changes**.

e) Click **Next**.

Step 8 Click **Next**.

Step 9 Deploy the changes to the registered managed devices.

The event data is not logged to the SAL (OnPrem) until the logging policy changes are deployed to the registered Firewall Threat Defense devices.

Note
If you must change any of these configurations, run the wizard again. If you disable the configuration or run he wizard again, all settings except the account credentials are retained.

You can view and work with these remotely stored events in the event viewer and context explorer in the management center, and include them when generating reports. You can also cross-launch from an event in the management center to view related data on your Secure Network Analytics Manager.

For more information, see the online help for the management center.

About Security Analytics and Logging(SaaS) for Firewall Threat Defense

SAL (SaaS) allows you to capture connection, intrusion, file, malware, and Security Intelligence events from all of your Firewall Threat Defense devices and view them in one place in Security Cloud Control. The events are stored in the Cisco cloud and are viewable from the **Event Logging** page in Security Cloud Control, where you can filter and review them to gain a clear understanding of what security rules are triggering in your network.

Licensing for SAL (SaaS)

The SAL (SaaS) licenses allows you to use a Security Cloud Control tenant to view firewall logs from all your managed devices at one place. For details on the available SAL (SaaS) licensing options, see the Cisco Security Analytics and Logging Ordering Guide.

Configure the SAL (SaaS) Integration

To deploy this integration, you must set up event data storage in SAL (SaaS) using either syslog or a direct connection.

- Send Cloud-Delivered Firewall Management Center-Managed Events to SAL (SaaS) Using Syslog, on page 553
- Send Cloud-Delivered Firewall Management Center-Managed Event Logs to SAL (SaaS) Using a Direct Connection, on page 555

Requirements and Guidelines

Type	Description
Cisco Secure Firewall Threat Defense	• Security Cloud Control-managed standalone Firewall Threat Defense devices, Version, 7.2 and later. • To send events using syslog, you must have Firewall Threat Defense device, Version 6.4 or later. • To send events directly, you must have Firewall Threat Defense device, Version 7.2 or later. • To optionally exclude individual Firewall Threat Defense devices from sending events directly, you must have Firewall Threat Defense device, Version 7.4.1 or later. • Your firewall system must be deployed and successfully generating events.
Regional cloud	• Determine the regional cloud that you want to send events to. • Events cannot be viewed from or moved between different regional clouds. • If you use a direct connection to send events to the Cisco Security Cloud for integration with Cisco SecureX, or Cisco SecureX threat response, or Cisco XDR, you must use the same cloud region for this integration. • If you send events directly, the regional cloud you specify in Security Cloud Control must match the region of your Security Cloud Control tenant.
Security Analytics and Logging subscription plan	• You must have a valid Security Analytics and Logging subscription plan. For more information about the subscription plans, see Security Analytics and Logging Licenses. • If your do not have a Security Analytics and Logging plan subscribed, you can request a 90-day trial by logging in to Security Cloud Control and navigating to **Events & Logs** > **Events** > **Event Logging** tab. You can purchase the desired subscription plan by following the instructions in the Security Cloud Control Firewall Management Ordering Guide. • You must buy a data plan that reflects the number of events the Cisco cloud receives from your threat defense devices daily. This is called your daily ingest rate. • Use the Logging Volume Estimator Tool to estimate your data storage requirements.

Type	Description
Accounts	When you purchase a license for this integration, you are provided with a Security Cloud Control tenant account to support the integration.

Type	Description
Connectivity	

Type	Description
	The Firewall Threat Defense devices must be able to connect outbound on port 443 to the Cisco Security Cloud at the following addresses: • US region: • api-sse.cisco.com • mx*.sse.itd.cisco.com • dex.sse.itd.cisco.com • eventing-ingest.sse.itd.cisco.com • registration.us.sse.itd.cisco.com • us.manage.security.cisco.com • edge.us.cdo.cisco.com • EU region: • api.eu.sse.itd.cisco.com • mx*.eu.sse.itd.cisco.com • dex.eu.sse.itd.cisco.com • eventing-ingest.eu.sse.itd.cisco.com • registration.eu.sse.itd.cisco.com • eu.manage.security.cisco.com • edge.eu.cdo.cisco.com • Asia (APJ) region: • api.apj.sse.itd.cisco.com • mx*.apj.sse.itd.cisco.com • dex.apj.sse.itd.cisco.com • eventing-ingest.apj.sse.itd.cisco.com • registration.apj.sse.itd.cisco.com • apj.cdo.cisco.com • edge.apj.cdo.cisco.com • Australia region: • api.aus.sse.itd.cisco.com • mx*.aus.sse.itd.cisco.com • dex.au.sse.itd.cisco.com

Type	Description
	• eventing-ingest.aus.sse.itd.cisco.com • registration.au.sse.itd.cisco.com • aus.cdo.cisco.com • India region: • api.in.sse.itd.cisco.com • mx*.in.sse.itd.cisco.com • dex.in.sse.itd.cisco.com • eventing-ingest.in.sse.itd.cisco.com • registration.in.sse.itd.cisco.com • in.cdo.cisco.com

Send Cloud-Delivered Firewall Management Center-Managed Events to SAL (SaaS) Using Syslog

This procedure provides information about the configuration for sending syslog messages for security events (connection, security intelligence, intrusion, file, and malware events) from devices managed by Security Cloud Control.

Before you begin

- Configure policies to generate security events, and verify that the events you expect to see are displayed in the applicable tables under the **Events & Logs** menu.
- Gather information relating to the syslog server IP address, port, and protocol (UDP or TCP).
- Ensure that your devices can reach the syslog server.

Procedure

Step 1 In the left pane, click **Administration** > **Firewall Management Center** to open the **Services** page.

Step 2 Click and select **Cloud-Delivered FMC** and then click **Configuration**.

Step 3 Configure the syslog settings for your threat defense device:

a) Click **Devices** > **Platform Settings** and edit the platform settings policy that is associated with your threat defense device.

b) In the left-side navigation pane, click **Syslog** and configure the syslog settings as follows:

Click this UI Element...	To Do the Following:
Logging Setup	Enable logging, specify FTP server settings, and the Flash usage.
Logging Destination	Enable logging to specific destinations and to specify filtering by message severity level, event class, or by a custom event list.
E-mail Setup	Specify the email address that is used as the source address for syslog messages that are sent as emails.
Events Lists	Define a custom event list that includes an event class, a severity level, and an event ID.
Rate Limit	Specify the volume of messages being sent to all the configured destinations and define the message severity level to which you want to assign the rate limits.
Syslog Settings	Specify the logging facility, enable the inclusion of a time stamp, and enable other settings to set up a server as a syslog destination.
Syslog Servers	Specify the IP address, protocol that is used, format, and security zone for the syslog server that is designated as a logging destination.

c) Click **Save**.

Step 4 Configure the general logging settings for the access control policy (including file and malware logging):

a) Click **Policies** > **Access Control** and then edit the access control policy that is associated with your threat defense device.

b) Click **More** and then choose **Logging**. Configure the general logging settings for the access control policy (including file and malware logging) as follows:

Click this UI Element...	To Do the Following:
Send using specific syslog alert	Select a syslog alert from the list of existing predefined alerts or add one by specifying the name, logging host, port, facility, and severity.
Use the syslog settings configured in the FTD Platform Settings policy deployed on the device	Unify the syslog configuration by configuring it in **Platform Settings** and reuse the settings in the access control policy. The selected severity is applied to all the connection and intrusion events. The default severity is **ALERT**.
Send Syslog messages for IPS events	Send events as syslog messages. The default syslog settings are used unless you override them.

Click this UI Element...	To Do the Following:
Send Syslog messages for File and Malware events	Send file and malware events as syslog messages. The default syslog settings are used unless you override them.

c) Click **Save**.

Step 5 Enable logging for security intelligence events for the access control policy:

a) In the same access control policy, click the **Security Intelligence** tab.

b) Click the logging icon and enable security intelligence logging using the following criteria:

- By Domain Name—Click the logging icon next to the **DNS Policy** drop-down list.
- By IP address—Click the logging icon next to **Networks**.
- By URL—Click the logging icon next to **URLs**.

c) Click **Save**.

Step 6 Enable syslog logging for each rule in the access control policy:

a) In the same access control policy, click the **Access Control** tab.

b) Click a rule to edit.

c) Click the **Logging** tab in the rule.

d) Check the **Log at beginning of connection** and **Log at end of connection** check boxes.

e) If you want to log file events, check the **Log Files** check box.

f) Check the **Syslog Server** check box.

g) Verify that the rule is **Using default syslog configuration in Access Control Logging**.

h) Click **Confirm**.

i) Click **Apply** to save the rule.

j) Repeat steps 7.a through 7.h for each rule in the policy and click **Save** to save the policy.

Step 7 If you have made all the required changes, deploy your changes to the managed devices.

Send Cloud-Delivered Firewall Management Center-Managed Event Logs to SAL (SaaS) Using a Direct Connection

Configure the Cloud-Delivered Firewall Management Center to send events directly to SAL (SaaS). Follow this procedure to enable the Cisco cloud event global setting in the Cloud-Delivered Firewall Management Center. When needed, you can exclude individual Firewall Threat Defense devices from sending event logs to SAL (SaaS). For more information, see Enable or Disable Threat Defense Devices to Send Event logs to SAL (SaaS) Using a Direct Connection.

Before you begin

- Onboard devices to the Cloud-Delivered Firewall Management Center, assign licenses to these devices, and configure these devices to send events directly to SAL (SaaS).

- Enable connection logging on a per-rule basis by editing a rule and choosing the **Log at Beginning of Connection** and **Log at End of Connection** options.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 In the left pane, click **Administration** > **Firewall Management Center**.

Step 3 Click **Cloud-Delivered FMC**, and in the **System** pane that is located at the right-side, click **Cisco Cloud Events**.

Step 4 In the **Configure Cisco Cloud Events** widget, do the following:

a. Click the **Send Events to the Cisco Cloud** toggle button to enable the overall configuration.

b. Check the **Send Intrusion Events to the cloud** check box to send the intrusion events to the cloud.

Note
When you enable sending intrusion events, the Firewall Threat Defense device sends the captured packet data to the cloud along with the intrusion events. Cisco Talos uses this packet data to analyze and eliminate false threats. This ensures focus on actual threats, enhancing the accuracy and efficiency of Talos' threat recommendations.

c. Check the **Send File and Malware Events to the cloud** check box to send the file and malware events to the cloud.

d. Choose an option to send the connection events to the cloud:

- Click the **None** radio button to not send connection events to the cloud.
- Click the **Security Events** radio button to send only security intelligence events to the cloud.
- Click the **All** radio button to send all the connection events to the cloud.

e. Click **Save**.

Enable Individual Threat Defense Devices to Send Events to SAL (SaaS) Using a Direct Connection

Enable or disable the FTD devices managed by the Cloud-Delivered Firewall Management Center to send events directly to SAL (SaaS). This device-level control allows you to optionally exclude specific FTD devices from sending event logs to the Cisco cloud to reduce traffic or to maintain a combination of SAL and on-premises event log storage.

Note

- To enable or disable sending events to the Cisco cloud from the FTD devices, enable the Cisco cloud event global setting in the Cloud-Delivered Firewall Management Center. For more information on enabling the Cisco cloud event global setting, see Send Cloud-Delivered Firewall Management Center-Managed Event Logs to SAL (SaaS) Using a Direct Connection, on page 555.

  Sending events to the Cisco cloud is enabled by default for all FTD devices when the Cisco cloud event global setting is enabled in the Cloud-Delivered Firewall Management Center.

- The option to enable or disable FTD devices to send event logs to the cloud is supported on FTD Version 7.4.1 or later.

Before you begin

- Onboard devices to the Cloud-Delivered Firewall Management Center, assign licenses to these devices, and configure these devices to send events directly to SAL (SaaS).
- Enable connection logging on a per-rule basis by editing a rule and choosing the **Log at Beginning of Connection** and **Log at End of Connection** options.

Procedure

Step 1 Log in to Security Cloud Control.
Step 2 From the left pane, click **Security Devices**.
Step 3 Click the **Devices** tab to view the device.
Step 4 Click the **FTD** tab to view FTD devices.
Step 5 Choose the FTD devices whose configurations you want to edit, from the **Security Devices** list.
Step 6 In the **Device Management** pane, click **Cloud Events**.
Step 7 Click the **Send Events to the Cisco Cloud** toggle button to enable or disable the configuration.
Step 8 Click **Save**.

View and Work with Events in Security Cloud Control

Procedure

Step 1 Log in to Security Cloud Control.
Step 2 From the Security Cloud Control menu, choose **Events & Logs** > **Events** > **Event Logging**.
Step 3 Use the **Historical** tab to view all historical events data. By default, the viewer displays this tab.
Step 4 To view the live events, click the **Live** tab.

What to do next

For more information about what you can do on this page, see the Security Cloud Control online help.

View AI Defense Events in Security Cloud Control

AI Defense events record activities within your AI environment. They capture interactions such as prompts submitted, responses generated, and any triggered AI Defense policy violations. Each event includes key information, such as the application and AI model involved, the specific rule triggered, and the action taken.

Before you begin

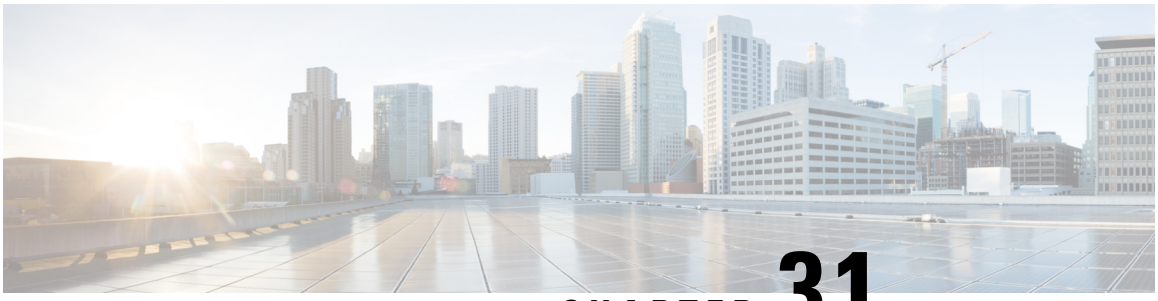
- Enable the AI Defense service in your Security Cloud Control tenant. For information on enabling the AI Defense service, see the AI Defense User Guide.
- You must be managing Firewall Threat Defense devices with Version 10.0 or higher.
- Enable your Firewall Threat Defense device to send events to the cloud. For more information, see Send Cloud-Delivered Firewall Management Center-Managed Event Logs to SAL (SaaS) Using a Direct Connection, on page 555.

Procedure

Step 1 In the navigation pane, choose **Events & Logs** > **Events** > **Event Logging**.

Step 2 Click the filter icon ( ).

Step 3 Check the **AI Defense** check box.



CHAPTER 31

FTD Dashboard

This chapter provides information about the FTD dashboard and describes the predefined widgets that offer an overview of your system.

- About the FTD Dashboard, on page 559
- View the FTD Dashboard, on page 560
- FTD Dashboard Widgets, on page 561
- Modify Time Settings for the FTD Dashboard, on page 564

About the FTD Dashboard

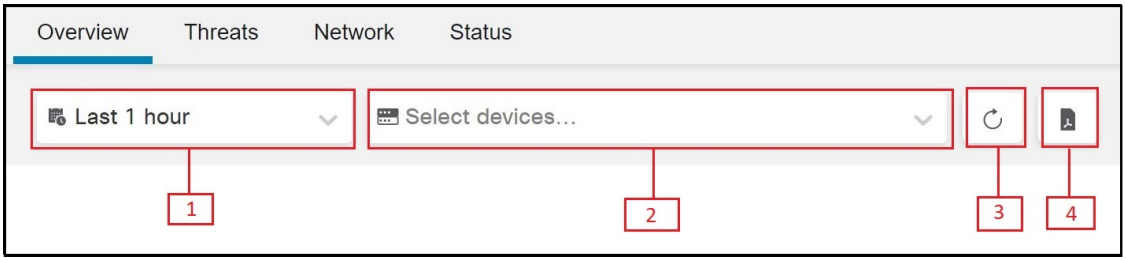
The FTD dashboard provides you an at-a-glance view of the status, including events data collected and generated by all Security Cloud Control-managed Firewall Threat Defense devices.

You can use this dashboard to view collective information that is related to the device status and the overall health of the devices in your deployment. The information that the FTD dashboard provides depends on how you license, configure, and deploy the devices in your system. Although, the FTD dashboard displays data for all Security Cloud Control-managed Firewall Threat Defense devices, you can choose to filter device-based data. You can also choose the time range to display for specific time range.

This dashboard uses tabs to display predefined widgets: small, self-contained components that provide insights into different aspects of the system. For example, the Network Activity widget shows you event graphs that display information about all connection, malware, and intrusion events. The widgets in the dashboard are predefined and cannot be customized. This dashboard is visible to all Security Cloud Control users who have access to a Security Cloud Control tenant.

- The dashboard does not displays any event statistics for historic events.
- Since the aggregation service batch processes the events to aggregate every five minutes, you can expect a five-minute latency between the time the events are aggregated to the time the statistics are displayed.

Figure 36: FTD Dashboard



Number	Description
1	Allows you to change the time range to reflect a period as short as the last hour or as long as the last year. When you change the time range, the widgets automatically update the events data to reflect the new time range.
2	Allows you to filter the events data based on selected devices. If no devices are selected, the widgets display all available events data.
3	Reinitiates the events data query
4	Displays the events data in a PDF output format. You can choose to download or save a copy of this PDF on your local machine.

View the FTD Dashboard

From the Security Cloud Control menu, choose **Insights & Reports** > **Firewall Threat Defense** > **Dashboard** to view the **FTD Dashboard**.

By default, the home page for your tenant displays the **Overview** tab.

The dashboard includes widgets that are listed under each tab: Threat, Network, Applications, and Status tab.

The following table lists the widgets available under each tab:

Name of the Tab	Available Widgets
Overview	All available widgets
Threat	• Top Intrusion Rules • Top Intrusion Attackers • Top Intrusion Targets • Top Malware Signatures • Top Malware Senders • Top Malware Receivers • Malware Events by Disposition

Name of the Tab	Available Widgets
Network	• Network Activity • Event Activity • Access Control Actions • Top Access Control Policies • Top Access Control Rules • Top Devices • Top Users • Top Users by Blocked Connections
Status	• Top Devices with Health Alerts • Top Loaded Devices
Applications	• Top Web Applications • Top Client Applications • Blocked Web Applications

FTD Dashboard Widgets

The FTD dashboard displays predefined widgets that can provide you with at-a-glance views of the current system status. These views include:

- Data about the events collected and generated by Firewall Threat Defense devices managed FMC.
- Information about the status and the overall health of the devices in your deployment.

Top Intrusion Rules Widget

The **Top Intrusion Rules** widget displays counts of the intrusion events that have occurred over the specified time range and are organized by priority. These counts include statistics on intrusion events with dropped packets and different impacts. The generated list is scrollable.

Top Intrusion Attackers Widget

The **Top Intrusion Attackers** widget displays counts of intrusion events for the top attacking host IP addresses (causing those events) on your monitored network.

Top Intrusion Targets Widget

The **Top Intrusion Targets** widget displays counts of intrusion events for the top target host IP addresses (targeted in the connections causing those events) on your monitored network.

Top Malware Signatures Widget

The **Top Malware Signatures** widget displays counts of the top malware signatures detected in the network traffic for top file-sending host IP addresses.

Top Malware Senders Widget

The **Top Malware Senders** widget displays counts of the top malware threats that are detected in the network traffic for top file-sending host IP addresses.

Top Malware Receivers Widget

The **Top Malware Receivers** widget displays counts of the top malware threats that are detected in the network traffic for all top file-receiving host IP addresses.

Malware Events by Disposition Widget

The **Malware Events by Disposition** widget displays counts of all disposition of malware events that are generated when the managed device detects a file containing malware.

Network Activity Widget

The **Network Activity** widget displays all ingress and egress data rate that is based on information from the connection events.

Event Activity Widget

The **Event Activity** widget displays counts of events that have occurred in the last hour and the total number of each event type that is available in the database.

Access Control Actions Widget

The **Access Control Actions** widget displays counts of events that are logged based on the allowed or blocked access control actions for each event. If you hover on the pie chart, you can view the percentage of the allowed and blocked actions.

Top Access Control Policies Widget

The **Top Access Control Policies** widget displays counts of top access control policies generating events.

Top Access Control Rules Widget

The **Top Access Control Rules** widget displays the top five counts of the access control rules that are used for each event. These counts can be sorted by bytes or events.

Top Devices Widget

The **Top Devices** widget displays counts of events per devices. These count can be sorted by bytes or events.

Top Users Widget

The **Top Users** widget displays a list of users on your monitored network that are associated with the highest intrusion event counts. It draws data primarily from the intrusion detection (IDS) User Statistics and Intrusion Events tables. It displays authoritative user data.

Top Users by Blocked Connections Widget

The **Top Users by Blocked Connections** widget displays users who have the highest number of blocked connection events associated with them in your monitored network.

Top Devices with Health Alerts Widget

The **Top Devices with Health Alerts** widget displays the current compiled health status for Security Cloud Control-managed Firewall Threat Defense devices.

Top Loaded Devices Widget

The **Top Loaded Devices** widget displays a list of Secure Firewall Threat Defense devices along with the CPU usage information.

Top Web Applications Widget

The **Top Web Applications** widget displays the web applications that are associated with the highest number of connection events in your monitored network. You can sort this list based on bytes or event count.

Top Client Applications Widget

The **Top Client Applications** widget displays the most used client applications in your monitored network. You can sort this list based on bytes or event count.

Top Blocked Web Applications Widget

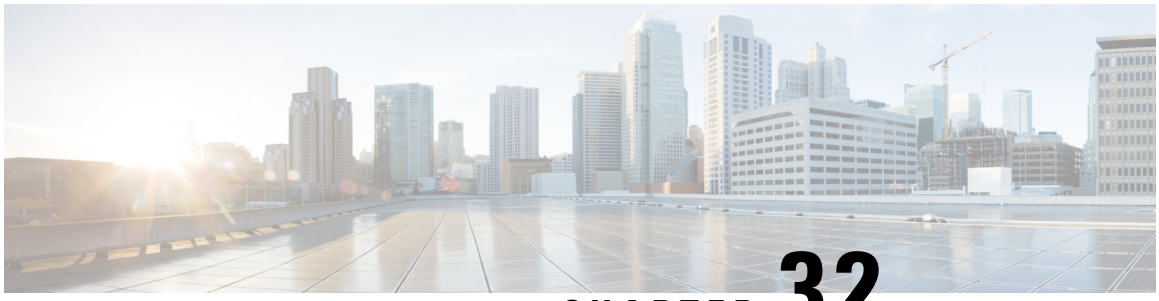
The **Top Web Applications** widget displays the top web applications that are associated with blocked connection events and the count of blocked connection events in your monitored network.

Modify Time Settings for the FTD Dashboard

You can change the time range to reflect a period as short as the last hour (the default) or as long as the last year. When you change the time range, the widgets that can be constrained by time automatically update to reflect the new time range.

The maximum number of data points in any graph is 300, and the time setting determines how much time is summarized within each data point. Following is the number of data points, and the time span covered, in the FTD dashboard for each time range:

- 1 hour = 12 data points, 5 minutes each
- 6 hours = 72 data points, 5 minutes each
- 1 day = 288 data points, 5 minutes each
- 1 week = 300 data points, 33.6 minutes each
- 2 weeks = 300 data points, 67.2 minutes each
- 30 days = 300 data points, 144 minutes each
- 90 days = 300 data points, 432 minutes each
- 180 days = 300 data points, 864 minutes each
- 1 year = 300 data points, 1752 minutes each



CHAPTER 32

Connection Logging

The following topics describe how to configure the system to log connections made by hosts on your monitored network:

- Connection logging, on page 565
- Connection logging limitations, on page 574
- Best practices for connection logging, on page 574
- Requirements and prerequisites for connection logging, on page 576
- Connection logging, on page 576
- Application-aware and protocol-aware syslogs, on page 581

Connection logging

Connection logging is a system feature that

- generates logs of connections detected by managed devices known as *connection events*,
- provides granular control through rules and policies over which connections to log, when to log them, and where to store the data, and
- creates special connection events called *security-related connection events* that represent connections blocked by the reputation-based Security Intelligence feature.

Connection event data

Connection events contain data about the detected sessions. The specific data available for any event can vary, but generally includes:

- Basic connection properties: timestamp, source and destination IP address, ingress and egress zones, the device that handled the connection, and so on
- Additional connection properties discovered or inferred by the system: applications, requested URLs, or users associated with the connection, and so on
- Metadata about why the connection was logged: which configuration handled the traffic, whether the connection was allowed or blocked, details about encrypted and decrypted connections, and so on

Log connections according to the security and compliance needs of your organization. When setting up connection logging, remember that the system can log a connection for multiple reasons. Disabling logging in one place does not guarantee that matching connections will not get logged elsewhere.

The information contained in a connection event depends on factors such as traffic characteristics and relevant configuration.



Note You can supplement the connection logs gathered by your managed devices with connection data generated from exported NetFlow records. This is particularly useful if you have NetFlow-enabled routers or other devices monitoring networks that managed devices cannot access.

Logged connections

Logged connections are network connections that

- are automatically saved to the Cloud-Delivered Firewall Management Center database as end-of-connection events,
- generate logs regardless of other logging configurations unless connection event storage is disabled, and
- include connections associated with intrusions, file and malware events, intelligent application bypass, and monitored traffic.

Connection types and logging behavior

The system automatically logs specific connection types based on their association with security events and traffic monitoring.

Connection types that generate automatic logs:

- Connections associated with intrusions: The system automatically logs connections associated with intrusion events, unless the connection is handled by the access control policy's default action.
- Connections associated with file and malware events: The system automatically logs connections associated with file and malware events.
- Connections associated with intelligent application bypass: The system automatically logs bypassed and would-have-bypassed connections associated with IAB.
- Monitored connections: The system always logs the ends of connections for monitored traffic, even if the traffic matches no other rules and you do not enable default action logging.



Note File events generated by inspecting NetBIOS-SSN (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.

Special behavior for default action connections:

- When an intrusion policy associated with the access control default action generates an intrusion event, the system does *not* automatically log the end of the associated connection. Instead, you must explicitly enable default action connection logging.
- However, if you enable beginning-of-connection logging for the default action, the system *does* log the end of the connection when an associated intrusion policy triggers, in addition to logging the beginning of the connection.

For more information about monitored connections, refer to Logging for monitored connections, on page 568.

Other Connections You Can Log

So that you log only critical connections, enable connection logging on a per-rule basis. If you enable connection logging for a rule, the system logs all connections handled by that rule.

You can also log connections handled by policy default actions. Depending on the rule or default action (and for access control, a rule's inspection configuration), your logging options differ.

Prefilter Policy: Rules and Default Action

You can log connections (including entire plaintext, passthrough tunnels) that you fastpath or block with a prefilter policy.

Prefiltering uses outer-header criteria to handle traffic. For tunnels that you log, the resulting connection events contain information from the outer, encapsulation headers.

For traffic subject to further analysis, logging in the prefilter policy is disabled, although matching connections may still be logged by other configurations. The system performs all further analysis using inner headers, that is, the system independently handles and logs each connection within an allowed tunnel.

Decryption Policy: Rules and Default Action

You can log connections that match a decryption rule or decryption policy default action.

For blocked connections, the system immediately ends the session and generates an event. For monitored connections and connections that you pass to access control rules, the system generates an event when the session ends.

Access Control Policy: Security Intelligence Decisions

You can log a connection whenever it is blocked by the reputation-based Security Intelligence feature.

Optionally, and recommended in passive deployments, you can use a monitor-only setting for Security Intelligence filtering. This allows the system to further analyze connections that would have been blocked by Security Intelligence, but still log the match. Security Intelligence monitoring also allows you to create traffic profiles using Security Intelligence information.

When the system logs a connection event as the result of Security Intelligence filtering, it also logs a matching Security Intelligence event, which is a special kind of connection event that you can view and analyze separately, and that is also stored and pruned separately. So that you can identify the matching IP address in the connection, host icons beside blocked and monitored IP addresses look slightly different in the tables on the pages under the **Analysis** > **Connections** > **Events** menus.

Access Control Policy: Rules and Default Action

You can log connections that match an access control rule or access control policy default action.

Related Topics

Rules and policy actions effect on logging, on page 568

Rules and policy actions effect on logging

Rules and policy actions effect on logging is a traffic handling mechanism that

- determines how the system inspects and handles matching traffic,
- controls when and how details about matching traffic can be logged, and
- provides metadata about why the connection was logged, including which configurations handled the traffic.

Connection event metadata

Connection events contain metadata about why the connection was logged, including which configurations handled the traffic. Where you can configure connection logging, rule actions, and policy default actions determine not only how the system inspects and handles matching traffic, but also when and how you can log details about matching traffic.

Logging for fastpathed connections

Logging for fastpathed connections is a monitoring capability that

- records connection events for traffic that matches tunnel rules and prefilter rules with **Fastpath** actions,
- captures information for non-encrypted tunnels including outer session details for tunnel rules, and
- provides limited information in connection events since fastpathed traffic bypasses access control and QoS processing.

Supported traffic types

You can log fastpathed connections and non-encrypted tunnels, which include traffic matching these rules and actions in the prefilter policy:

- Tunnel rules—**Fastpath** action (logs the outer session)
- Prefilter rules—**Fastpath** action

Fastpathed traffic bypasses the rest of access control and QoS, so connection events for fastpathed connections contain limited information.

Logging for monitored connections

Logging for monitored connections is a system function that

- automatically logs connection ends for traffic matching specific monitoring configurations,
- operates independently of default action logging settings, and

- consolidates multiple Monitor rule matches into single connection events.

Monitored traffic configurations

The system always logs the ends of connections for traffic matching these configurations, even if the traffic matches no other rules and you do not enable default action logging:

- Security Intelligence: Block lists set to monitor (also generates a Security Intelligence event)
- SSL rules: **Monitor** action
- Access control rules: **Monitor** action

The system does not generate a separate event each time a single connection matches a Monitor rule. Because a single connection can match multiple Monitor rules, each connection event can include and display information on the first eight Monitor access control rules that the connection matches, as well as the first matching SSL Monitor rule.

Similarly, if you send connection events to an external syslog or SNMP trap server, the system does not send a separate alert each time a single connection matches a Monitor rule. Rather, the alert that the system sends at the end of the connection contains information on the Monitor rules the connection matched.

Logging for trusted connections

Logging for trusted connections is a system capability that

- records the beginnings and ends of trusted connections for traffic matching access control rules with **Trust** action
- captures events for traffic matching the access control default action **Trust All Traffic**, and
- generates limited information because trusted connections are not subject to deep inspection or discovery.

Trusted connection logging behavior

You can log the beginnings and ends of trusted connections, which include traffic matching these rules and actions:

- Access control rules: **Trust** action
- Access control default action: **Trust All Traffic**



Note Although you *can* log trusted connections, we recommend you do not do so because trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

TCP connections detected by a Trust rule on the first packet generate only an end-of-connection event. The system generates the event one hour after the final session packet.

Logging for blocked connections

Logging for blocked connections is a security feature that

- records traffic matching blocking rules and actions across multiple policy types,
- generates connection events with different timing based on the blocking mechanism, and
- provides visibility into denied traffic for security monitoring and analysis.

Blocked connection types

You can log blocked connections, which include traffic matching these rules and actions:

- Tunnel rules—**Block**
- Prefilter rules—**Block**
- Prefilter default action—**Block all tunnel traffic**
- Security Intelligence—Block lists not set to Monitor (also generates a Security Intelligence event)
- Decryption rules—**Block** and **Block with reset**
- SSL default action—**Block** and **Block with reset**
- Access control rules—**Block**, **Block with reset**, and **Interactive Block**
- Access control default action—**Block All Traffic**

Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.



Caution Firewall Threat Defense device logs end-of-connection events for TCP connections where the three-way handshake does not complete, including SYN-only, SYN-SYN/ACK (no ACK), and SYN-RST/ACK scenarios (except when TLS Identity is enabled). These logs provide visibility into incomplete or suspicious connection attempts. However, logging blocked TCP connections during a Denial of Service (DoS) attack can affect system performance and overwhelm the database with multiple similar events. Before you enable logging for a Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.

Beginning vs. end-of-connection logging for blocked connections

When you log a blocked connection, how the system logs it depends on why the connection was blocked. This is important to keep in mind when configuring correlation rules based on connection logs:

- For decryption rules and decryption policy default actions that block encrypted traffic, the system logs **end**-of-connection events. This is because the system cannot determine if a connection is encrypted using the first packet in the session.
- For other blocking actions, the system logs **beginning**-of-connection events. Matching traffic is denied without further inspection.

Logging bypassed interactive blocks

Interactive blocking access control rules, which cause the system to display a warning page when a user browses to a prohibited website, allow you to configure end-of-connection logging. This is because if the user

clicks through the warning page, the connection is considered a new, allowed connection which the system can monitor and log.

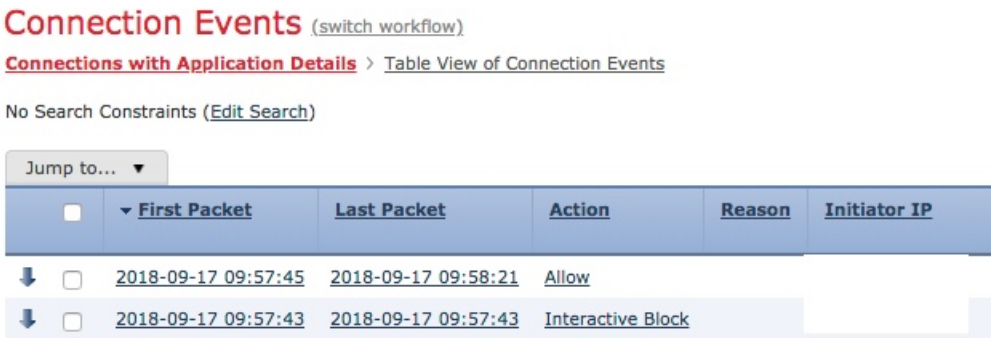
Therefore, for packets that match an Interactive Block or Interactive Block with Reset rule, the system can generate these connection events:

- A beginning-of-connection event when a user's request is initially blocked and the warning page is displayed; this event has an associated action of **Interactive Block** or **Interactive Block with Reset**
- Multiple beginning- or end-of-connection events if the user clicks through the warning page and loads the originally requested page; these events have an associated action of **Allow** and a reason of **User Bypass**

Interactive block example

The following figure shows an example of an interactive block followed by allow.

Figure 37: Interactive block followed by allow



Logging for allowed connections

Logging for allowed connections is a security feature that

- generates file and malware events when file policies detect or block files during traffic inspection, and
- provides configurable logging options on a per-access-control-rule basis.

Connection logging details

You can log allowed connections, which includes traffic matching the following rules and actions:

- SSL rules—**Decrypt** action
- SSL rules—**Do not decrypt** action
- SSL default action—**Do not decrypt**
- Access control rules—**Allow** action
- Access control default action—**Network Discovery Only** and any intrusion prevention option

Enabling logging for these configurations ensures the connection is logged, while also permitting (or specifying) the next phase of inspection and traffic handling. SSL logging is always end-of-connection; access control configurations also allow beginning-of-connection logging.

Although the **Analyze** action in tunnel and prefilter rules also allows connections to continue with access control, logging is disabled for rules with this action. Matching connections may still be logged by other configurations. Allowed tunnels might have their encapsulated sessions evaluated and logged individually.

When you allow traffic with an access control rule or default action, you can use an associated intrusion policy to further inspect traffic and block intrusions. For access control rules, you can also use a file policy to detect and block prohibited files, including malware. Unless you disable connection event storage, the system automatically logs most allowed connections associated with intrusion, file, and malware events. For detailed information, refer to Logged connections, on page 566.

Connections with encrypted payloads are not subject to deep inspection, so connection events for encrypted connections contain limited information.

When a file policy detects or blocks a file, it logs one of these events to the Cloud-Delivered Firewall Management Center database:

- *File events*, which represent detected or blocked files, including malware files.
- *Malware events*, which represent detected or blocked malware files only.
- *Retrospective malware events*, which are generated when the malware disposition for a previously detected file changes.

You can disable this logging on a per-access-control-rule basis. You can also disable file and malware event storage entirely.



Note We recommend you leave file and malware event logging enabled.

Connection logging modes

A connection logging mode is a connection monitoring approach that

- provides two distinct timing options for capturing connection events,
- handles blocked traffic differently than allowed traffic, and
- requires optimization choices to balance performance with information completeness.

Connection logging timing and exceptions

You can log a connection at its beginning or its end, with specific exceptions for blocked traffic:

- Blocked traffic: Because blocked traffic is immediately denied without further inspection, only beginning-of-connection events are typically available for logging; there is no unique end-of-connection to log.
- Blocked encrypted traffic: When you enable connection logging in a decryption policy, the system logs end-of-connection rather than beginning-of-connection events. This occurs because the system cannot

determine if a connection is encrypted with the first packet and therefore cannot immediately block encrypted sessions.

To optimize performance, log either the beginning or the end of any connection, but not both. Monitoring a connection for any reason forces end-of-connection logging. For a single non-blocked connection, the end-of-connection event contains all information present in the beginning-of-connection event, as well as data gathered over the duration of the session.

The table presents differences between beginning and end-of-connection events, along with the advantages of logging each.

Table 39: Comparing beginning and end-of-connection events

Attributes	Beginning-of-Connection Events	End-of-Connection Events
Can be generated...	When the system detects the beginning of a connection (or, after the first few packets if event generation depends on application or URL identification).	When the system: • Detects the close of a connection. • Does not detect the end of a connection after a period of time. • Can no longer track the session due to memory constraints.
Can be logged for...	All connections except those blocked by the decryption policy.	Most connections.
Contain...	Only information that can be determined in the first packet (or the first few packets, if event generation depends on application or URL identification).	All information in the beginning-of-connection event, plus information determined by examining traffic over the duration of the session; for example, the total amount of data transmitted or the timestamp of the last packet in the connection. **Note** The connection event does not count the amount of data transmitted after the threat defense returns a snort verdict for the connection or if you fastpath the connection.
Are useful...	If you want to log: • Blocked connections. • Only the beginning of a connection because the end-of-connection information does not matter to you.	If you want to: • Log encrypted connections handled by a decryption policy. • Perform any kind of detailed analysis on, or trigger correlation rules using, information collected over the duration of the session. • View connection summaries (aggregated connection data) in custom workflows, view connection data in graphical format, or create and use traffic profiles.

Connection logging limitations

Connection logging limitations are restrictions that prevent certain types of network connections from being logged by firewall systems.

Connection types that cannot be logged

Connection logging cannot capture information for these connection types:

- The outer session of a plaintext, passthrough tunnel whose encapsulated connections are inspected by access control.
- TCP connections if the three-way handshake is not completed, to avoid denial-of-service attacks against your firewalls. To monitor or debug failed connections, you can use the packet capture feature (Packet capture, on page 440).

Best practices for connection logging

Best practices for connection logging are guidelines that

- ensure you log only critical connections by enabling connection logging on a per-access-control-rule basis,
- help avoid logging unnecessary traffic such as trusted connections or infrastructure traffic not related to your service offering, and
- optimize storage and performance by preventing duplicate logging and configuring appropriate event limits.

Connections that are always logged

The system automatically logs these connections:

- Some connections associated with detected files, malware, intrusions, and Intelligent Application Bypass (IAB).

  For more information, refer to Logged connections, on page 566.

- Monitored connections.

  For more information, refer to Logging for monitored connections, on page 568.

Connections to never log

Do not enable logging for the following:

- Access control rules with a Trust action.

  Trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

- Do not enable logging for Block rules in passive deployments. To log connections that the system would have blocked if your devices were deployed inline, use a Monitor rule instead of a Block rule.

Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

- Traffic you're not interested in. Examples follow:
  - Specific allowed traffic, such as DNS requests to a trusted DNS host.
  - Infrastructure traffic that is not related to your service offering.

(As previously mentioned, you can still monitor this traffic for threats.)

As discussed in Logged connections, on page 566, even if you disable logging for the preceding, intrusion events, malware, and IAB are still logged.

Avoid logging what is being logged elsewhere

If another device or service is logging connection data for a network segment, disable logging for that segment's data in the Cloud-Delivered Firewall Management Center. Examples:

- If a router logs connection events on the same network segment as the Cloud-Delivered Firewall Management Center, avoid logging the same connections on the Cloud-Delivered Firewall Management Center unless you need those connection events for something else, such as correlation policies or traffic profiles.
- If you use Secure Network Analytics to leverage NetFlow records reported from switches and routers to identify potential behavioral anomalies and suspicious traffic patterns, you can disable connection logging for rules monitoring those segments and instead rely on Secure Network Analytics for behavioral analytics for those parts of your network.

  For more information, consult the Secure Network Analytics documentation.

Log either the beginning or end of the connection (not both)

If you have a choice between beginning and end-of-connection logging, enable end-of-connection logging. This is because end-of-connection logs information from beginning-of-connection events, as well as information gathered over the duration of the session.

Log the beginning of connections only if you want to log blocked connections, or if end-of-connection information does not matter to you.

For more information, refer to Connection logging modes, on page 572.

Logging for blocked traffic

Because blocked traffic is immediately denied without further inspection, usually you can log only beginning-of-connection events.

For more information, refer to Logging for blocked connections, on page 569.

Log events to an external location

If your company's security policy permits it, you can save disk space on your Cloud-Delivered Firewall Management Center by streaming logs to an external source using any of the following:

- eStreamer, which enables you to stream logs from a Cloud-Delivered Firewall Management Center to a custom-developed client application. For more information, refer to the *Secure Firewall Management Center Event Streamer Integration Guide* .
- Syslog or SNMP trap are referred to as alert responses. For more information, refer to Configuring external alerts with alert responses, on page 487.

Control what is displayed in connection events

To specify the number of rows displayed in connection events, click your username in the upper right of the Cloud-Delivered Firewall Management Center, click **User Preferences**, and then click the **Event View Settings**tab. The maximum you can set is 1000 events per page.

Requirements and prerequisites for connection logging

This reference provides the technical requirements and prerequisites that must be met to implement and use connection logging features.

Model support

Any.

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Connection logging

Connection logging is a network monitoring feature that records network connection activities to match various rules and conditions.

Logging connections with tunnel and prefilter rules

This task enables logging for prefilter policy rules to monitor and analyze connection events for troubleshooting and security analysis purposes.

The prefilter policy applies to Secure Firewall Threat Defense devices only.

Before you begin

- Set the rule action to **Block** or **Fastpath**. Logging is disabled for the **Analyze** action, which allows connections to continue with access control, where other configurations determine their handling and logging.
- Logging is performed on inner flows, not on the encapsulating flow.

Follow these steps to configure logging for tunnel and prefilter rules:

Procedure

Step 1 In the prefilter policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click **Logging**.

Step 3 Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.

To optimize performance, log either the beginning or the end of any connection, but not both. Because blocked traffic is immediately denied without further inspection, you cannot log end-of-connection events for Block rules.

Step 4 Specify where to send connection events:

Step 5 Click **Save** to save the rule.

Step 6 Click **Save** to save the prefilter policy.

What to do next

Deploy configuration changes.

Configure logging for decryptable connections with TLS/SSLDecryption rules

Logging decryptable connections allows you to track and analyze traffic that is being monitored by decryption rules. This is particularly important for security analysis and compliance requirements.

Procedure

Step 1 In the decryption policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click **Logging**.

Step 3 Check the **Log at End of Connection** check box.

For monitored traffic, end-of-connection logging is required.

Step 4 Specify where to send connection events.

Send events to the event viewer if you want to perform Cloud-Delivered Firewall Management Center-based analysis on these connection events. For monitored traffic, this is required.

Step 5 Click **Save** to save the rule.

Step 6 Click **Save** to save the decryption policy.

What to do next

Deploy configuration changes.

Logging connections with security intelligence

Enable Security Intelligence logging to monitor and analyze network connections, IP addresses, URLs, and domain names for security threats and policy compliance.

The Security Intelligence policy requires the Threat Smart License or Protection Classic License.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**.

Step 2 Click the **Logging** ( ) icon to enable Security Intelligence logging using the following criteria:

- By IP address: Click the logging icon next to **Networks**.
- By URL: Click the logging icon next to **URLs**.
- By Domain Name: Click the logging icon next to the **DNS Policy** drop-down list.

If the logging icon is disabled, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Check the **Log Connections** check box.

Step 4 Specify where to send connection and Security-Related connection events.

Send events to the Security Cloud Control if you want to perform Security Cloud Control-based analysis, or if you set a Block list to monitor-only.

Step 5 Click **OK** to set logging options.

Step 6 Click **Save** to save the policy.

What to do next

Deploy configuration changes.

Logging connections with access control rules

Depending on your choices for the rule action and deep inspection options, your logging options differ; refer to Rules and policy actions effect on logging, on page 568.

Procedure

Step 1 In the access control policy editor, click **Edit** ( ) next to the rule where you want to configure logging.

If **View** ( ) appears instead, the configuration is inherited from an ancestor policy, belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Click **Logging**.

Step 3 Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.

To optimize performance, log either the beginning or the end of any connection, but not both.

Step 4 (Optional) Check the **Log Files** check box to log file and malware events associated with the connection.

It is recommended to leave this option enabled.

Step 5 Specify where to send the connection events:

- **Event Viewers**: Send events to the Cloud-Delivered Firewall Management Center. When using cloud management, send events to the cloud-delivered Cloud-Delivered Firewall Management Center and to an on-premises Cloud-Delivered Firewall Management Center if you have configured it to perform event analytics only. You can view the events in the event viewer of either product.
- **Syslog Server**: Send connection events to the syslog server configured in the Logging tab in Access Control Policy, unless overridden.

  Show Overrides: Displays the options to override the settings configured in the access control policy.

  - **Override Severity**: When you choose this option and select a severity for the rule, connection events for this rule will have the selected severity regardless of the severity configured in the Logging tab in Access Control Policy.
  - **Override Default Syslog Destination**: Send the syslog generated for the connection event for this rule to destination specified in this alert.

- **SNMP Trap**: Connection events are sent to the selected SNMP trap.

Step 6 Click **Confirm**.

Step 7 Click **Apply** to save the rule.

What to do next

Deploy configuration changes.

Logging connections with a policy default action

A policy's default action determines how the system handles traffic that matches none of the rules in the policy (except Monitor rules in access control and decryption policies, which match and log—but do not handle or inspect—traffic).

Logging settings for the decryption policies default action also govern how the system logs undecryptable sessions.

Before you begin

For prefilter default action logging, set the default action to **Block all tunnel traffic**. Logging is disabled for the **Allow all tunnel traffic** action, which allows connections to continue with access control, where other configurations determine their handling and logging.

Procedure

Step 1 In the policy editor, click the **Default Logging and Inspection** ⚙ next to the **Default Action** drop-down list.

Step 2 Specify when you want to log matching connections:

- Log at Beginning of Connection: Not supported for SSL default actions.
- Log at End of Connection: Not supported if you choose the access control **Block All Traffic** default action or the prefilter **Block all tunnel traffic** default action.

To optimize performance, log either the beginning or the end of any connection, but not both.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration. In an access control policy, the configuration may also be inherited from an ancestor policy.

Step 3 Specify where to send connection events.

Send events to the event viewer if you want to perform Cloud-Delivered Firewall Management Center-based analysis on these connection events.

Step 4 Click **Apply**.

Step 5 Click **Save** to save the policy.

What to do next

Deploy configuration changes.

Limiting logging of long URLs

End-of-connection events for HTTP traffic record the URL requested by monitored hosts. Disabling or limiting the number of stored URL characters may improve system performance. Disabling URL logging (storing zero characters) does not affect URL filtering. The system filters traffic based on requested URLs even though the system does not record them.

Procedure

Step 1 In the access control policy editor, click **More** > **Advanced Settings**, then click **Edit** (✎) next to **General Settings**.

If **View** (👁) appears instead, the configuration is inherited from an ancestor policy, belongs to an ancestor domain, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Enter the **Maximum URL characters to store in connection events**.

Step 3 Click **OK**.

Step 4 Click **Save** to save the policy.

What to do next

Deploy configuration changes.

Application-aware and protocol-aware syslogs

Application-aware and protocol-aware syslogs are a logging mechanism that

- generate logs that include application-specific and protocol-specific traffic data,
- provide enhanced network visibility through comprehensive data collection, and
- use the Snort 3 engine's deep packet inspection capabilities to extract protocol data.

Advanced logging protocols

The Firewall Threat Defense device generates these logs alongside existing connection, intrusion, file, and malware events. As syslogs, you can send these logs to external event management solutions for analysis. When enabled, advanced logging provides a configurable solution for sending these logs from the Firewall Threat Defense device to destination such as Splunk or a syslog alert server.

Advanced logging supports these protocols.

Table 40: Protocols supported by advanced logging

Protocol	Description
CONN	Logs the data collected from end-of-connection events, including the transport protocol, service protocol, and session duration.
DNS	Logs the data collected from responses to TCP and UDP DNS sessions, including commands, their arguments, response codes, and response messages.

Protocol	Description
FTP	Logs the commands and data exchanged during an FTP session, including the command and its arguments, response codes, and response messages.
HTTP	Logs the data collected from the HTTP responses, including the HTTP method, header, and status code.
NOTICE	Logs the data collected for intrusion events, including the GID:SID pair, rule messages, and associated references.
WEIRD	Logs the data about anomalies detected in traffic flow.

For more information about application log fields, refer to Advanced Logging Syslog Fields, on page 587.

Advanced logging guidelines and limitations

Advanced logging guidelines and limitations are operational constraints and performance considerations that

- affect network performance when configured without proper filtering,
- require selective configuration to avoid device memory consumption issues, and
- have specific platform and version requirements for support.

Advanced logging restrictions and requirements

Consider these guidelines when implementing advanced logging:

- Advanced logging can degrade network performance if used without filters in the access control rule. To reduce the volume of logs, use access control rules to filter specific traffic types and use networks and ports to limit logging to a particular network configuration.
- Enabling advanced logging for all protocols can impact the Firewall Threat Defense device's performance. Configure advanced logging selectively to avoid increased device memory consumption.
- Advanced logging is supported only on Snort 3-based Firewall Threat Defense devices running Version 10.0 or later.
- Advanced logging does not support sending logs to a syslog server configured in the platform settings. To send events to the default logging destination configured in the access control policy, you must configure a syslog alert server as the default logging destination.
- Advanced logging does not support Splunk profiles that use data interfaces for sending events to Splunk. To send events to Splunk, you must configure the Splunk profile to use the management interface for sending events. For more information about configuring a Splunk profile, refer to Splunk integration: Send events directly from Cloud-Delivered Firewall Management Center , on page 527.

Enable advanced logging

Configure the Firewall Threat Defense device to generate Snort 3 inspector logs that include application-specific and protocol-specific data and send these logs to event management solutions for analysis.

Caution Advanced logging might cause a performance drop within the network if used without the filters configured in the access control rule. Filter the specific traffic types using the access control rules to reduce the volume of logged traffic. Use the networks and ports in the access control rules to limit logging to a particular network configuration.

Before you begin

- Ensure that connection logging is enabled in your access control policy.
- Ensure that you have configured the logging destination, such as Splunk or the local syslog server, to which you want to send the logs.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** ( ) next to the access control policy that you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 In the access control policy editor, select **Advanced Settings** from the **More** drop-down menu at the end of the packet flow line.

Step 4 Click the **Edit** ( ) icon next to **Advanced Logging**.

Step 5 Check the **Enable advanced logging** check box to enable logging.

Step 6 Choose the log file format. By default, JSON format is selected.

Step 7 You can send event logs to the default logging destination configured in the access control policy's logging settings, or to one or more of the following destinations:

- **All Splunk profiles**
- **Syslog**

Note

- If you choose to send event logs to the default logging destination configured in the access control policy, note that advanced logging does not support sending logs to the syslog server configured in the platform settings. You must configure a syslog alert server as the default logging destination.
- If you choose the **All Splunk profiles** option, note that advanced logs are sent to all Splunk profiles configured for a Firewall Threat Defense device.
- Advanced logging does not support Splunk profiles that use data interfaces for sending events to Splunk. To send events to Splunk, you must configure the Splunk profiles to use the management interface for sending events.

Caution

Sending advanced logging syslog to multiple destinations might impact device's performance.

Step 8 Click **Save**.

Step 9 Click **Save** to save the policy.

What to do next

Deploy configuration changes.

Configure protocols for advanced logging

Advanced logging for application protocols allows you to monitor and analyze network connections by capturing detailed information about specific application protocols. This feature is configured at the access control policy level and can be applied to individual rules within the policy.

Before you begin

Ensure that you have enabled the advanced logging feature for the access control policy that you want to modify.

Follow these steps to configure protocols for advanced logging:

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** ( ) next to the access control policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 In the access control policy editor, you have the options to add a new rule or edit an existing rule.

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click **Edit** ( ).

Step 4 Click **Advanced Logging**.

Step 5 Check the check box next to the application protocol for which you want to log the connection.

Caution
Enabling logging for all the protocols in a Firewall Threat Defense device might affect its performance, because the device must extract data, convert it to targeted formatting, and send it to the configured destination.

Step 6 Click **Confirm**.

Step 7 Click **Apply** to save the rule.

Step 8 Click **Save** to save the policy.

What to do next

Deploy configuration changes.

Monitor and troubleshoot advanced logging

To collect statistics of the advanced logging events and to get health alerts, you must enable the **Snort 3 Statistics** module in your Firewall Threat Defense device's health policy (**System** (⚙) > **Health** > **Policy**). When enabled, you can monitor metrics such as the number of advanced logging events that are dropped, failed to send, or sent successfully.

To view these health metrics, navigate to **System** (⚙) > **Health** > **Monitor** and then add the metrics to a dashboard by selecting them from the **Snort 3 Performance Statistics** group.

Advanced Logging events failed to transmit to syslog servers

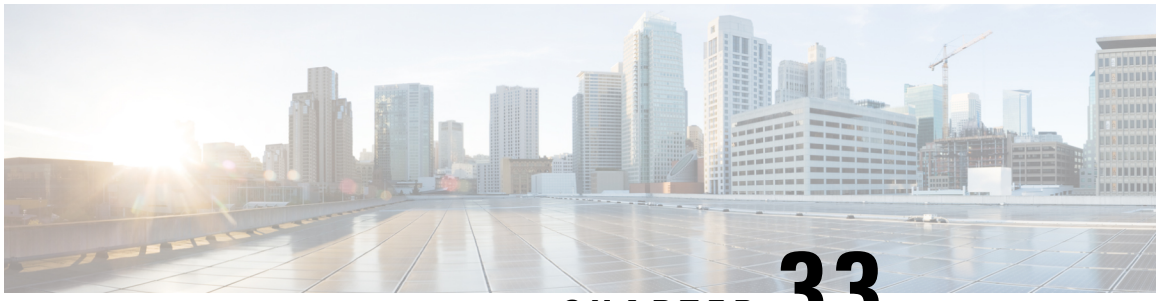
This alert appears when the syslog messages fail to transmit due to a connection issue with the syslog server or a configuration error.

- Check the status of your syslog server to ensure it is operational and accessible.
- Verify the syslog configuration within the Cloud-Delivered Firewall Management Center for any errors.

Advanced logging events to syslog servers were dropped

This alert appears when syslog messages are dropped due to memory overflow in the Firewall Threat Defense device, suggesting that the device cannot process logs at the current rate.

Review your advanced logging configuration. Consider applying more granular filters in your access control rules to reduce the volume of logs generated, or selectively enable advanced logging for fewer protocols, to decrease device memory consumption. For more information, refer to Advanced logging guidelines and limitations, on page 582.



CHAPTER 33

Advanced Logging Syslog Fields

Advanced logs are generated and sent as syslog messages to the configured destinations for consumption by an external tool. The syslog fields for all advanced logging protocols are detailed here.

- Common fields, on page 587
- CONN protocol fields, on page 588
- DNS protocol fields, on page 589
- FTP fields, on page 590
- HTTP fields, on page 591
- Notice protocol fields, on page 592
- Weird protocol fields, on page 592

Common fields

These fields appear in the syslog message across all protocols:

Table 41: HTTP fields

Field	Description
id.orig_h	The client IP address involved in a connection.
id.orig_p	The client TCP or UDP port used for a connection.
id.resp_h	The server IP address involved in a connection.
id.resp_p	The server TCP or UDP port used for a connection.
pkt_num	The packet number within a network flow.
tenant_id	The identifier for a tenant associated with an event.
ts	The timestamp of the packet that triggered the log record shows when the event occurred.
uid	A unique connection ID that enables you to correlate log records related to the same network flow.

CONN protocol fields

Conn protocol fields capture various aspects of network connections, including state, duration, event history, byte counts, packet counts, transport protocol, and application service information.

Fields available in the CONN protocol for connection analysis:

Field	Description
conn_state	Captures the state of the connection based on the protocol in use. • **UDP**: States include CLT_SRV_UDP_SEEN (packets from both client and server observed), CLT_UDP_SEEN (only client packets observed), and SRV_UDP_SEEN (only server packets observed). • **TCP**: Tracks the client and server states independently using prefixes CLT_ (client) and SRV_ (server), reflecting the TCP STATE machine per RFC standards, with additional states for MID-STREAM activity (TCP_MID_STREAM_SENT, TCP_MID_STREAM_REC) and TCP_STATE_NONE. • **Other traffic**: Indicates non-UDP and non-TCP traffic or error cases.
duration	The duration of the connection, in seconds.

history	A code that indicates the event sequence of the connection.
	Each letter in the history code represents a specific event. Uppercase letters indicate client-side events, and lowercase letters indicate server-side events. Events are recorded only once per direction.
	• For UDP, events include: `d` (packet with payload).
	• For TCP, events include:
	  • `s` (SYN)
	  • `h` (SYN-ACK)
	  • `a` (pure ACK or PUSH)
	  • `d` (packet with payload)
	  • `f` (FIN)
	  • `r` (reset)

orig_bytes	The total number of TCP or UDP payload bytes transmitted by the client during the connection.
orig_pkts	The number of packets sent by the originator.
proto	The transport layer protocol of a connection, for example, IP, ICMP, TCP, or UDP.

Field	Description
resp_bytes	The total number of TCP or UDP payload bytes transmitted by the server during the connection.
resp_pkts	The number of packets sent by the responder.
service	A connection's application protocol. This value indicates the last detected service on the traffic flow.

DNS protocol fields

This reference provides detailed information about DNS protocol fields, including boolean values, resource records, transport protocols, and various identifiers used in DNS transactions.

DNS protocol fields contain specific data elements that define the behavior and content of DNS transactions:

Field	Description
AA	A boolean value that indicates whether this value is an authoritative answer (AA) to a query. For example, T or F.
addl	The list of additional responses. This list contains all the resource records (RR) found in the **addl** section. The resource record types are handled in the same way as for the **answers** field.
answers	A list of resource records that directly answer a DNS query. Each resource record provides specific information about a domain. This may include its IP address, mail server, or other properties, depending on the type of query. All resource records appear in the answers section of a DNS response. The decoding process represents each resource record by summarizing its contents. Each resource record type has specific decoding rules, depending on the type of information it represents. The following resource record types contains type-specific information when decoded—A, AAAA, BIND9 signing, CNAME, DNSKEY, DS, LOC, MX, NS, NSEC, OPT, PTR, RRSIG, SOA, SPF, SRV, SSHFP, TXT. All other resource record types are decoded by the default decoder. If the resource record type is not known or not specifically handled, it is displayed as UNKNOWN followed by the resource record type numeric value.
auth	The list of authoritative responses. This list contains all the resource records found in the **auth** section. The resource record types are handled the same way as those for the **answers** field.
proto	The transport protocol used for the DNS connection, TCP or UDP.
qclass	A 16-bit integer specifying the class of a DNS query.
qclass_name	A descriptive name for the class of a DNS query.
qtype	A 16-bit integer specifying the type of a DNS query.
qtype_name	A descriptive name for the type of a DNS query.

Field	Description
query	The domain name that is the subject of a DNS transaction.
RA	A boolean value indicating the availability of recursive query support in a server, for example, T or F.
rcode	A 16-bit integer specifying the response code to a DNS query.
rcode_name	A descriptive name for the response code to a DNS query.
RD	A boolean value indicating whether a client asked the server to pursue the query recursively, for example, T or F.
rejected	A boolean value indicating whether the server responded with an error code and no query, for example, T or F.
TC	A boolean value indicating whether a message was truncated because of UDP PDU size limits, for example, T or F.
trans_id	A 16-bit identifier assigned to a DNS query.
TTLs	The list of caching intervals for the corresponding answers. The values in the list are separated by an empty space.
Z	A 3-bit integer set to 0 unless Domain Name System Security Extensions (DNSSEC) is used.

Values in the list are separated by an empty space.

FTP fields

FTP protocol analysis relies on specific fields that capture essential command parameters, data channel information, file transfers, and server responses. Understanding these fields enables accurate interpretation of FTP activity and troubleshooting.

The following fields are commonly used during FTP protocol analysis:

Field	Description
arg	The parameters associated with an FTP command.
command	The last FTP command seen in a session.
data_channel.orig_h	The IP address of a data channel originator.
data_channel.passive	A boolean value indicating whether passive mode was used for a data channel, for example, true or false.
data_channel.resp_h	The IP address of a data channel receiving point.
data_channel.resp_p	The TCP port of a data channel receiving point.
file_size	The size of a file transferred during an FTP session.

Field	Description
reply_code	The FTP reply code from a server in response to a command.
reply_msg	The FTP reply message from a server.
user	The username used for an FTP session.

HTTP fields

This reference provides details about HTTP protocol fields that are captured and available in application event logs for analysis and monitoring purposes.

Field	Description
host	The host header from an HTTP request, indicating the target server.
info_code	The last informational status code returned by a server.
info_msg	The last informational reason phrase returned by a server.
method	The HTTP method used in a request, for example, GET or POST.
origin	The origin header from a client.
orig_filenames	List of file names sent by a client. Values in the list are separated by empty spaces.
orig_mime_types	List of the content type (MIME type) files sent by a client. Values in the list are separated by empty spaces.
proxied	List of headers associated with proxied requests. Values in a list are separated by empty spaces.
referrer	The referrer header, indicating the URL of a page that is linked to the requested resource.
request_body_len	The length of an HTTP request body (decompressed and normalized).
response_body_len	The length of an HTTP response body (decompressed and normalized).
resp_filenames	List of file names sent by a server. Values in the list are separated by empty spaces.
resp_mime_types	List of the content type files sent by a server. Values in the list are separated by empty spaces.
status_code	The HTTP status code returned by a server, for example, 200 or 404.
status_msg	The HTTP status message returned by a server, for example, OK or Not Found.
trans_depth	The number of request-response pairs seen in an HTTP session.

Field	Description
uri	The Uniform Resource Identifier (URI) from an HTTP request, specifying the resource being requested.
user_agent	The user-agent header from a client, identifying the client software.
version	The HTTP version used in a request.

Notice protocol fields

This reference provides definitions for the protocol fields that appear in notice logs when intrusion policy rules are triggered, helping you understand the information captured in security event logs.

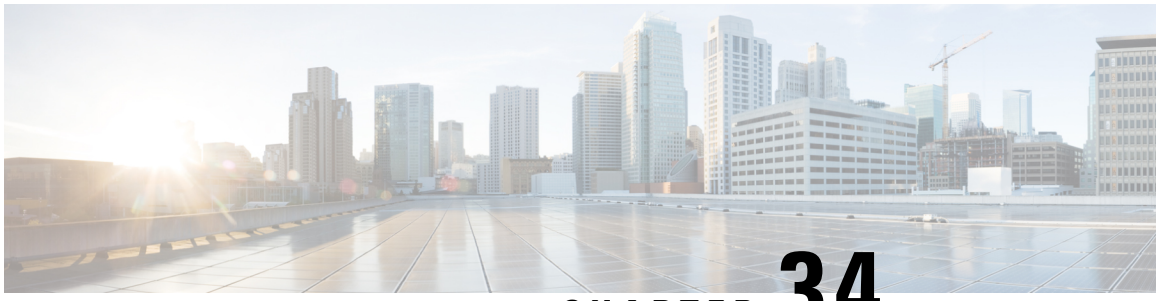
The notice protocol fields provide detailed information about intrusion events. These fields include:

Field	Description
action	The intrusion policy action that was configured for the triggered intrusion policy rule, for example alert, drop, or pass.
gid	The GID of the intrusion rule that triggered the log. This value is collected and displayed even if the rule is disabled.
msg	The message associated with the intrusion rule that triggered the log. This field provides a description of why the flow was logged.
proto	The transport layer protocol associated with the event, for example, IP, ICMP, TCP, or UDP.
rev	The revision number of the intrusion rule that was triggered.
refs	A list of references (URLs) associated with the intrusion rule. These references provide additional information about the specific threat or vulnerability the rule is designed to detect. The references are expanded to full URLs in the log.
sid	The SID of the intrusion rule that triggered the log. This value is collected and displayed even if the rule is disabled.
source	The name of the inspector module that was assigned to process the flow. This identifies the specific component within Snort that detected the anomaly. The name of the module appears similar to that in the Network Analysis Policy.

Weird protocol fields

This reference describes the various protocol fields that can appear in intrusion detection logs, including rule identifiers, messages, transport protocols, and inspector modules.

Field	Description
gid	The GID of the intrusion rule that triggered the log. This value is collected and displayed even if the rule is disabled.
msg	The message associated with the intrusion rule that triggered the log. This field provides a description of why the flow was logged.
proto	The transport layer protocol associated with the event, for example, IP, ICMP, TCP, or UDP.
sid	The SID of the intrusion rule that triggered the log. This value is collected and displayed even if the rule is disabled.
source	The name of the inspector module that was assigned to process the flow. This identifies the specific component within Snort that detected the anomaly. The name of the module appears similar to that in the Network Analysis Policy.



CHAPTER 34

User Identity and Activity

A identity source is a user data management system that

- collects active session data, user data, and user activity data from network users,
- organizes this data into distinct workflows for individual user analysis, and
- enables administrators to monitor and manage user behavior across the network.

Identity sources collect active session data, user data, and user activity data. The data is displayed in individual user-related workflows:

- Active Sessions: This workflow displays all current user sessions on your network. A single user running several simultaneous active sessions would occupy several rows in this table. For more information about the types of user data displayed in this workflow, refer to Active sessions data, on page 602.
- Users: This workflow displays all users seen on your network. A single user occupies a single row in this table. For more information about the types of user data displayed in this workflow, refer to User data, on page 603.
- User Activity: This workflow displays all user activity seen on your network. A single user with more than one instance of user activity would occupy several rows in this table. For more information about the types of user activity displayed in this workflow, refer to User activity data, on page 606.

- User-related fields, on page 595
- Active sessions data, on page 602
- User data, on page 603
- User activity data, on page 606
- User profile and host history, on page 608

User-related fields

User-related data is displayed in the active sessions, users, and user activity tables.

Note Active sessions for Azure AD realm users are displayed only in the **Active Sessions** new UI layout and not in the legacy UI.

Table 42: Active sessions, users, and user activity field descriptions

Field	Description	Active Sessions Table	Users Table	User Activity Table
Active Session Count	The number of active sessions associated with the user.	No	Yes	No
Authentication Type	The type of authentication: `No Authentication`, `Passive Authentication`, `Active Authentication`, `Guest Authentication`, `Failed Authentication`, or `VPN Authentication`. For more information about the supported identity sources for each Authentication Type, refer to the *Cisco Secure Firewall Management Center Device Configuration Guide* .	Yes	No	Yes
Available for Policy	A value of **Yes** means the user was retrieved from the user store (for example, Active Directory).) A value of **No** means the Cloud-Delivered Firewall Management Center received a report of a login for that user but the user is not in the user store. One way this can happen is if a user in an excluded group logs in to the user store. You can exclude groups from being downloaded when you configure a realm. Users not available for policy are recorded in the Cloud-Delivered Firewall Management Center but are not sent to managed devices.	No	Yes	No
Count	**Note** The **Count** field is displayed only after you apply a constraint that creates two or more identical rows. Depending on the table, the number of sessions, users, or activity events that match the information that appears in a particular row.	Yes	Yes	Yes
Current IP	(See also Current IP/Domain and IP address.) The IP address associated with the host that the user is logged into. This field is blank in the Users table if there are no active sessions for a user.	Yes	No	No

Field	Description	Active Sessions Table	Users Table	User Activity Table
Department	The user's department, as obtained by a realm. If there is no department explicitly associated with the user on your servers, the department is listed as whatever default group the server assigns. For example, on Active Directory, this is `Users (ad)`. This field is blank if: • You have not configured a realm. • The Cloud-Delivered Firewall Management Center cannot correlate the user in the Cloud-Delivered Firewall Management Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login).	Yes	Yes	No
Description	More information, if available, about the session, user, or user activity.	No	No	Yes
Device	For user activity detected by traffic-based detection or an active authentication identity source, the name of the device that identified the user. For other types of user activity, the managing Cloud-Delivered Firewall Management Center. **Note** If you have configured your VPN in a high-availability deployment, the device name displayed against active VPN sessions can be the primary or secondary device that identified the user session.	Yes	No	Yes
Discovery Application	The application or protocol used to detect the user. • For user activity detected by traffic-based detection, one of the following: ldap, pop3, imap, oracle, sip, http, ftp, mdns, or aim. **Note** Users are not added to the database based on SMTP logins. • For all other user activity: ldap.	Yes	Yes	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
Current IP Domain / **Domain**	In the Active Sessions table, the multitenancy domain where the user activity was detected. In the Users table, the multitenancy domain associated with the user's realm. In the User Activity table, the multitenancy domain where the user activity was detected. This field is only present if you have ever configured the Cloud-Delivered Firewall Management Center for multitenancy.	Yes	Yes	Yes
Email	The user's email address. This field is blank if: • The user was added to the database via an AIM login. • The user was added to the database via an LDAP login and there is no email address associated with the user on your LDAP servers.	Yes	Yes (as E-Mail)	No
End Port	If the user was reported by the TS Agent and their session is currently active, this field identifies the end value for the port range assigned to the user. This field is blank if the user's TS Agent session is inactive or if the user was reported by another identity source.	No	No	Yes
Endpoint Location	The IP address of the network device that used ISE to authenticate the user, as identified by ISE. If you do not configure ISE, this field is blank.	No	No	Yes
Endpoint Profile	The user's endpoint device type, as identified by Cisco ISE. If you do not configure ISE, this field is blank.	No	No	Yes
Event	The user activity event type.	No	No	Yes
First Name	The user's first name, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Cloud-Delivered Firewall Management Center cannot correlate the user in the Cloud-Delivered Firewall Management Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no first name associated with the user on your servers.	Yes	Yes	No

Field	Description	Active Sessions Table	Users Table	User Activity Table
IP Address	For **User Login** user activity, the IP address or internal IP address involved in the login: • LDAP, POP3, IMAP, FTP, HTTP, MDNS, and AIM logins — the address of the user's host • SMTP and Oracle logins — the address of the server • SIP logins — the address of the session originator (See also Current IP and Current IP/Domain.) An associated IP address does not mean the user is the current user for that IP address; when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. For other types of user activity, this field is blank.	No	No	Yes
Last Name	The user's last name, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Cloud-Delivered Firewall Management Center cannot correlate the user in the Cloud-Delivered Firewall Management Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no last name associated with the user on your servers.	Yes	Yes	No
Last Seen	The date and time that a session was last initiated (or user data was updated) for the user.	Yes	Yes	No
Login Time	The date and time that the session was initiated for the user.	Yes	No	No

Field	Description	Active Sessions Table	Users Table	User Activity Table
Phone Number	The user's telephone number, as obtained by a realm. This field is blank if: • You have not configured a realm. • The Cloud-Delivered Firewall Management Center cannot correlate the user in the Cloud-Delivered Firewall Management Center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no telephone number associated with the user on your servers.	Yes (as Phone)	Yes	No
Realm	The identity realm associated with the user.	Yes	Yes	Yes
Security Group Tag	The Security Group Tag (SGT) attribute applied by Cisco TrustSec as the packet entered a trusted TrustSec network. If you do not configure ISE, this field is blank.	No	No	Yes
Session Duration	The duration of the user session, calculated from the **Login Time** and the current time.	Yes	No	No
Start Port	If the user was reported by the TS Agent and their session is currently active, this field identifies the start value for the port range assigned to the user. This field is blank if the user's TS Agent session is inactive or if the user was reported by another identity source.	No	No	Yes
Time	The time that the system detected the user activity.	No	No	Yes
User	At minimum, this field displays the user's realm and username. For example, `Lobby\jsmith`, where Lobby is the realm and jsmith is the username. If a realm downloads additional user data from an LDAP server and the system associates it with a user, this field also displays the user's first name, last name, and type. For example, `John Smith (Lobby\jsmith, LDAP)`, where John Smith is the user's name and LDAP is the type. **Note** Because traffic-based detection can record unsuccessful AIM logins, the Cloud-Delivered Firewall Management Center may store invalid AIM users (for example, if the username is misspelled by the user).	Yes	Yes	No
Username	The username associated with the user.	Yes	Yes	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
VPN Bytes In	For Remote Access VPN-reported user activity, the total number of bytes received from the remote peer or client by the Firewall Threat Defense. **Note** You can view the total number of bytes received once the user's VPN session is terminated. For ongoing VPN sessions, this is not a dynamic counter. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Bytes Out	For Remote Access VPN-reported user activity, the total number of bytes transmitted to the remote peer or client by the Firewall Threat Defense. **Note** You can view the total number of bytes transmitted once the user's VPN session is terminated. For ongoing VPN sessions, this is not a dynamic counter. For other types of user activity, this field is blank.	No	No	Yes
VPN Client Application	For Remote Access VPN-reported user activity, the remote user's AnyConnect VPN module of Cisco Secure Client application. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Client Country	For Remote Access VPN-reported user activity, the country name as reported by the Secure Client VPN. For other types of user activity, this field is blank.	No	No	Yes
VPN Client OS	For Remote Access VPN-reported user activity, the remote user's endpoint operating system as reported by the Secure Client VPN. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Client Public IP	For Remote Access VPN-reported user activity, the publicly routable IP address of the Secure Client VPN device. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Connection Duration	For Remote Access VPN-reported user activity, the total time (HH:MM:SS) that the session was active. For other types of user activity, this field is blank.	No	No	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
VPN Connection Profile	For Remote Access VPN-reported user activity, the name of the connection profile (tunnel group) used by the VPN session. Connection profiles are part of a Remote Access VPN Policy. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Group Policy	For Remote Access VPN-reported user activity, the name of the group policy assigned to the client when the VPN session is established; either the statically-assigned group policy associated with the VPN Connection Profile, or the dynamically-assigned group policy if RADIUS is used for authentication. If assigned by the RADIUS server, this group policy overrides the static policy configured for the VPN Connection Profile. Group policies configure common attributes for groups of users in Remote Access VPN policies. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Session Type	For Remote Access VPN-reported user activity, the type of session: LAN-to-LAN or Remote. For other types of user activity, this field is blank.	Yes	No	Yes

Active sessions data

The **Integrations** > **Users** > **Active Sessions** workflow displays select information about current user sessions.

When a user on your network runs several sessions simultaneously, the system can uniquely identify the sessions if:

- they have unique **IP Address** values.
- they have unique **Start Port** and **End Port** values, as provided by the Cisco Terminal Services (TS) Agent.
- they have unique **Current IP Domain** values.
- they were authenticated by different identity sources.
- they were associated with different identity realms.

For more information about the user and user activity data stored by the system, refer to User data, on page 603 and User activity data, on page 606.

View active session data

You can view a table of active sessions, and then manipulate the event view depending on the information you are looking for.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

Procedure

Step 1 Choose **Analysis** > **Users heading** > **Active Sessions**.

Step 2 You have these options:

- Perform basic workflow actions.
- Learn more about the contents of the columns in the table; refer to Active sessions data, on page 602 and User-related fields, on page 595.

User data

This reference describes how user records are created and updated in the database and what identity-related information the system stores.

When an identity source reports a user login for a user who is not already in the database, the user is added to the database, unless you have specifically restricted that login type.

The system updates the users database when one of the following occurs:

- A user on the Cloud-Delivered Firewall Management Center manually deletes a non-authoritative user from the Users table.
- An identity source reports a logoff by that user.
- A realm ends the user session as specified by the realm's **User Session Timeout: Authenticated Users**, **User Session Timeout: Failed Authentication Users**, or **User Session Timeout: Guest Users** setting.



Note If you have ISE/ISE-PIC configured, you may see host data in the users table. Because host detection by ISE/ISE-PIC is not fully supported, you cannot perform user control using ISE-reported host data.

The type of user login that the system detected determines what information is stored about the new user.

Identity Source	Login Type	User Data Stored
ISE/ISE-PIC	Active Directory LDAP RADIUS RSA	• username • current IP address • Security Group Tag (SGT) — not supported with ISE-PIC • endpoint profile/device type — not supported with ISE-PIC • endpoint location/location IP — not supported with ISE-PIC • type (LDAP)
TS Agent	Active Directory	• username • current IP address • start port • end port • type (LDAP)
captive portal	Active Directory LDAP	• username • current IP address • type (LDAP)
traffic-based detection	LDAP AIM Oracle SIP HTTP FTP MDNS	• username • current IP address • type (AD)
	POP3 IMAP	• username • current IP address • email address • type (`pop3` or `imap`)

Note No data about Microsoft Azure Active Directory users is displayed in this table.

If you configure a realm to automatically download users, the Cloud-Delivered Firewall Management Center queries the servers based on the interval you specified. It may take five to ten minutes for the Cloud-Delivered Firewall Management Center database to update with user metadata after the system detects a new user login. The Cloud-Delivered Firewall Management Center obtains the following information and metadata about each user:

- username
- first and last names
- email address
- department
- telephone number
- current IP address
- Security Group Tag (SGT), if available
- endpoint profile, if available
- endpoint location, if available
- start port, if available
- end port, if available

The number of users the Cloud-Delivered Firewall Management Center can store in its database depends on your Cloud-Delivered Firewall Management Center model. When a non-authoritative user login is detected on a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user login is detected for that host, only another authoritative user login changes the current user.

Note that traffic-based detection of AIM, Oracle, and SIP logins create duplicate user records because they are not associated with any of the user metadata that the system obtains from LDAP servers. To prevent overuse of user count because of duplicate user records from these protocols, configure traffic-based detection to ignore those protocols.

You can search, view, and delete users from the database; you can also purge all users from the database.

For information about general user-related event troubleshooting, refer to *Cisco Secure Firewall Management Center Device Configuration Guide*.

View user data

You can view a table of users, and then manipulate the event view depending on the information you are looking for.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

Procedure

Step 1 Choose **Analysis** > **Users heading** > **Users**.

Step 2 Choose from the following options:

- Perform basic workflow actions.
- Learn more about the contents of the columns in the table; refer to User-related fields, on page 595.

User activity data

A user activity data record is a security event type that

- logs the appearance and login of users on a network,
- enables correlation between user actions and other network events such as intrusions, and
- supports automated remediation and alerting through customizable rules.

The system generates events that communicate the details of user activity on your network. When the system detects user activity, the user activity data is logged to the database. You can view, search, and delete user activity; you can also purge all user activity from the database.

The system logs a user activity event when a user is seen on your network for the first time. Subsequent appearances by that user do not log new user activity events. However, if the user's IP address changes, the system logs a new user activity event.

The system also correlates user activity with other types of events. For example, intrusion events can tell you the users who were logged into the source and destination hosts at the time of the event. This correlation can tell you who was logged into the host that was targeted by an attack, or who initiated an internal attack or portscan.

You can also use user activity in correlation rules. Based on the type of user activity as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and alert responses when network traffic meets your criteria.

Note If you have ISE/ISE-PIC configured, you may see host data in the users table. Because host detection by ISE/ISE-PIC is not fully supported, you cannot perform user control using ISE-reported host data.

Descriptions of the four types of user activity data follow.

New User Identity

This type of event is generated when the system detects a login by an unknown user that is not in the database.

The system logs a user activity event when a user is seen on your network for the first time. Subsequent appearances by that user do not log new user activity events. However, if the user's IP address changes, the system logs a new user activity event.

User Login

This type of event is generated when any of the following occur:

- Captive portal performs a successful or failed user authentication.
- Traffic-based detection detects a successful or failed user login.



Note SMTP logins detected by traffic-based detection are not recorded unless there is already a user with a matching email address in the database.

When a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user.

If you are using captive portal or traffic-based detection, note the following about failed user login and failed user authentication data:

- Failed logins reported by traffic-based detection (LDAP, IMAP, FTP, and POP3 traffic) are displayed in the table view of user activity, but not in the table view of users. If a known user failed to log in, the system identifies them by their username. If an unknown user failed to log in, the system uses **Failed Authentication** as their username.
- Failed authentications reported by captive portal are displayed in both the table view of user activity and the table view of users. If a known user failed to authenticate, the system identifies them by their username. If an unknown user failed to authenticate, the system identifies them by the username they entered.

Delete User Identity

This type of event is generated when you manually delete a user from the database.

User Identity Dropped: User Limit Reached

This type of event is generated when the system detects a user that is not in the database, but cannot add the user because you have reached the maximum number of users in the database as determined by your Cloud-Delivered Firewall Management Center model.

After you reach the user limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

However, the system favors authoritative users. If you have reached the limit and the system detects a login for a previously undetected authoritative user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new authoritative user.

User Indications of Compromise Events

The following user IOC changes are logged in the user activity database:

- When indications of compromise are resolved.
- When indication of compromise rules are enabled or disabled for users.

For information about general user-related event troubleshooting, refer to the *Cisco Secure Firewall Management Center Device Configuration Guide*.

View user activity data

You can view a table of user activity, and then manipulate the event view depending on the information you are looking for. The page you see when you access user activity differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of user activity and terminates in a user details page, which contains user details for every user that meets your constraints.

Procedure

Step 1 Choose **Analysis** > **Users heading** > **User Activity**.

Step 2 You have the following options:

- Perform basic workflow actions.
- Learn more about the contents of the columns in the table; refer to User-related fields, on page 595.

User profile and host history

A user profile is a web interface display that

- provides detailed information about a specific user through the User pop-up window (titled "User Identity" in the web interface)
- shows the same user data as you would see in the table view of users, and
- includes host history that provides a graphic representation of the last twenty-four hours of the user's activity.

User profile access points

You can learn more about a specific user by viewing the User pop-up window. The page that appears, called the "User Profile" in this document, is titled "User Identity" in the web interface.

You can display the window from:

- any event view that associates user data with other kinds of events
- the table view of active sessions
- the table view of users

User information also appears in the terminating page for users workflows.

The user data you see is the same as you would see in the table view of users.

Host history section

The host history provides a graphic representation of the last twenty-four hours of the user's activity. A list of IP addresses of the hosts that the user logged into and logged off of approximates login and logout times with bar graphs. A typical user might log on to and off of multiple hosts in the course of a day. For example, periodic automated logins to a mail server would display as multiple short sessions, while longer logins (such as during working hours) display longer sessions.

If you use traffic-based detection or captive portal to capture failed logins, the host history also includes hosts where the user failed to log in.

The data used to generate the host history is stored in the user history database, which by default stores 10 million user login events. If you do not see any data in the host history for a particular user, either that user is inactive, or you may need to increase the database limit.

View user details and host history

User details and host history provide valuable insights for security analysis and incident investigation.

Procedure

You have two options:

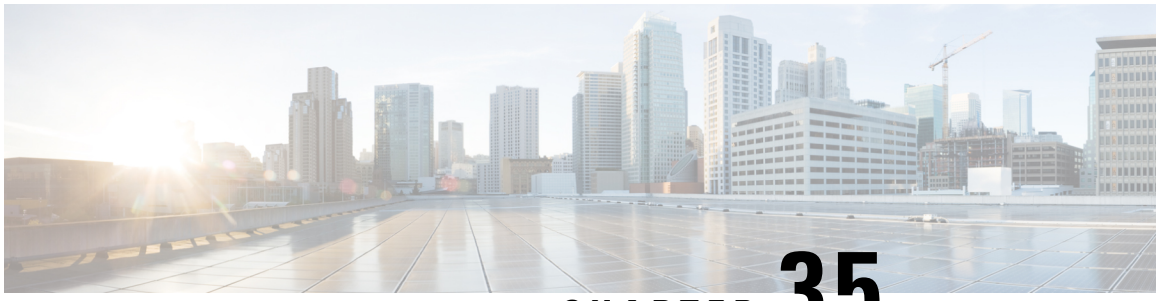
- In any event view that lists users, click user that appears next to a user identity **User icon**, or, for users associated with an indication of compromise, **Red User icon**.
- In any users workflow, click the Users terminating page.



PART XI

High Availability and Scalability

- Multi-Instance Mode, on page 613
- High Availability for Devices, on page 681
- Clustering: Secure Firewall 3100/4200, on page 723
- Clustering: Private Cloud , on page 799
- Clustering: Public Cloud, on page 845
- Clustering: Firepower 4100/9300, on page 935



CHAPTER 35

Multi-Instance Mode

You can deploy the Secure Firewall 3100/4200 as a single device (*appliance mode*) or as multiple container instances (*multi-instance mode*). This chapter describes how to deploy the device in multi-instance mode.

- Multi-instance mode, on page 613
- Licenses for instances, on page 625
- Requirements and prerequisites for instances, on page 625
- Guidelines and limitations for instances, on page 626
- Instance configuration, on page 629
- Multi-instance mode, on page 676

Multi-instance mode

Multi-instance mode is a deployment mode that enables multiple container instances on a single chassis that act as completely independent devices.

Multi-instance modes and appliance modes

A multi-instance mode and an appliance mode are deployment types for network devices that

- determine whether the device operates as a single instance or supports multiple, independent instances,
- offer distinct management and configuration options at both the chassis and instance levels, and
- provide different approaches to resource allocation, interface assignment, and monitoring.

You can run the device in either appliance mode or multi-instance mode.

Appliance mode

Appliance mode is the default. The device runs the native Firewall Threat Defense image and acts as a single device. The only chassis-level configuration available (on the **Chassis Manager** page) is for network module management (breakout ports or enabling/disabling a network module).

Multi-instance mode

If you change to multi-instance mode, the device runs the Secure Firewall eXtensible Operating System (FXOS) on the chassis, while each instance runs separate Firewall Threat Defense images. You can configure the mode using the FXOS CLI.

Because multiple instances run on the same chassis, you need to perform chassis-level management of:

- CPU and memory resources using resource profiles.
- Interface configuration and assignment.
- Deployment and monitoring of instances.

For a multi-instance device, you add the *chassis* to the Cloud-Delivered Firewall Management Center and configure chassis-level settings on the **Chassis Manager** page.

Chassis management interface

A chassis management interface is a network hardware feature that

- provides a dedicated pathway for device management and control,
- supports multi-instance modes by allowing each software instance its own management IP address, and
- separates management, eventing, and application access to streamline security, access, and monitoring.

Chassis management

The chassis uses the dedicated Management interface on the device. Multi-instance mode does not support using a data interface for chassis management or DHCP addressing for the Management interface.

You can only configure the chassis Management interface at the Firewall Threat Defense CLI (at initial setup) or FXOS CLI (after you convert to multi-instance mode). Refer to Change chassis management settings at the FXOS CLI, on page 673 to change Management interface settings in multi-instance mode.

Note By default, SSH is not allowed to this interface in multi-instance mode unless you enable the SSH server and an SSH access list. This difference means that you can connect to the application-mode threat defense Management interface using SSH, but after you convert to multi-instance mode, you can no longer connect using SSH by default. For more information, refer to Configure SSH and SSH access list, on page 654.

Instance management

All instances share the chassis management interface, and each instance has its own IP address on the Management network. After you add the instance and specify the IP address, you can make changes to the network settings at the Firewall Threat Defense CLI.

The instance Management IP address allows SSH by default.

Instance interfaces

A chassis interface is a network connection point that

- allows flexible assignment to instances through physical connections or VLAN subinterfaces,
- enables sharing (physical or VLAN) between multiple instances for efficient resource use, and
- restricts certain interface features (like subinterfaces or sharing) in multi-instance clusters, except for control links.

To provide flexible physical interface use for instances, you can create VLAN subinterfaces on the chassis and also share interfaces (VLAN or physical) between multiple instances. Refer to Shared interface scalability, on page 617 and Configure a subinterface, on page 640.



Note This chapter discusses *chassis* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense instance. For more information, refer to Chassis interfaces vs. instance interfaces, on page 615.

Interface types

An interface type is a network configuration element that

- determines how physical interfaces, VLAN subinterfaces, and EtherChannel interfaces are used in deployment, and
- defines whether interfaces can be shared among instances or dedicated to specific traffic.

Physical interfaces, VLAN subinterfaces, and EtherChannel interfaces can be one of these types:

- Data: Use for regular data or the failover link. Data interfaces cannot be shared between instances, and instances cannot communicate over the backplane to other instances. For traffic on Data interfaces, all traffic must exit the chassis on one interface and return on another interface to reach another instance. You can add VLAN subinterfaces to a data interface to provide separate failover links per High Availability pair.
- Data-sharing: Use for regular data. These data interfaces can be shared by one or more instances. Each instance can communicate over the backplane with all other instances that share this interface. Shared interfaces can affect the number of instances you can deploy. Shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode), inline sets, passive interfaces, or failover links.

Chassis interfaces vs. instance interfaces

Chassis interfaces vs. instance interfaces is an interface management model that

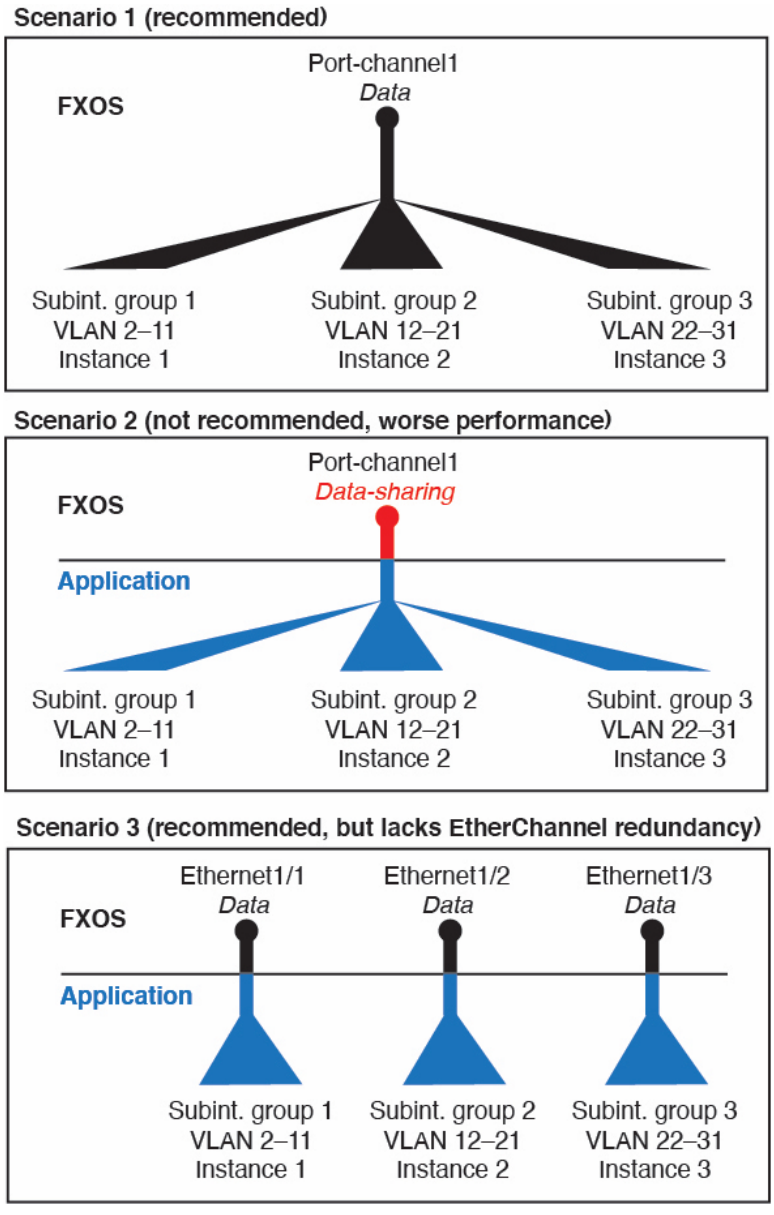
- manages basic Ethernet settings of physical interfaces, VLAN subinterfaces for instances, and EtherChannel interfaces at the chassis level,
- configures higher level settings within the instance, such as IP addresses for EtherChannels, and
- allows independent administrative control of interface states in both chassis and instance locations.

VLAN subinterfaces

You can create VLAN subinterfaces within the instance, just as you would for any device.

You can *also* create VLAN subinterfaces in the chassis. The instance-defined subinterfaces are not subject to the chassis limit. Choosing in which location to create subinterfaces depends on your network deployment and personal preference. For example, to share a subinterface, you must create the subinterface on the chassis. Another scenario that favors chassis subinterfaces comprises allocating separate subinterface groups on a single interface to multiple instances. For example, you want to use Port-channel1 with VLAN 2–11 on instance A, VLAN 12–21 on instance B, and VLAN 22–31 on instance C. If you create these subinterfaces in the instance, then you would have to share the parent interface in the chassis, which may not be desirable. See this illustration that shows the three ways you can accomplish this scenario:

Figure 38: VLANs in the chassis vs. the instance



Independent interface states in the chassis and in the instance

You can administratively enable and disable interfaces in both the chassis and in the instance. For an interface to be operational, the interface must be enabled in both locations. Because the interface state is controlled independently, you may have a mismatch between the chassis and instance.

The default state of an interface within the instance depends on the type of interface. For example, the physical interface or EtherChannel is disabled by default within the instance, but a subinterface is enabled by default.

Shared interface scalability

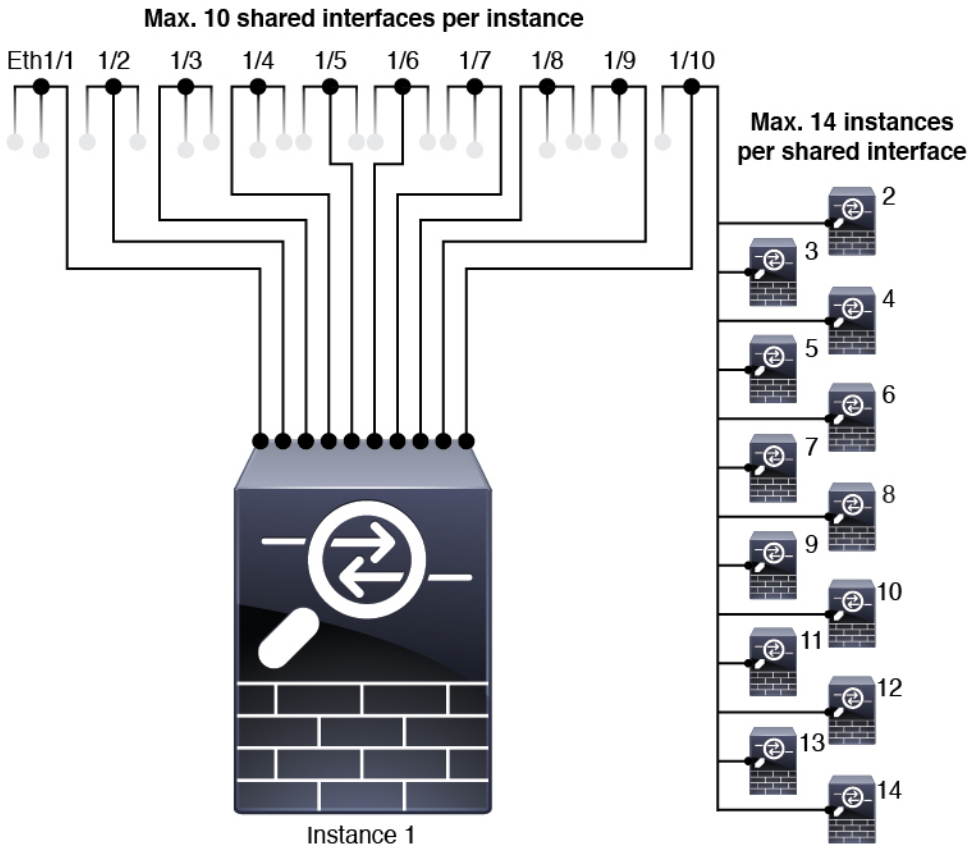
A shared interface is a network resource that

- allows multiple logical connections to transmit data concurrently,
- adapts bandwidth allocation based on real-time demand, and
- enhances scalability to support growing numbers of devices and traffic flows.

Instances can share data-sharing type interfaces. This capability lets you conserve physical interface usage as well as support flexible networking deployments. When you share an interface, the chassis uses unique MAC addresses to forward traffic to the correct instance. However, shared interfaces can cause the forwarding table to grow large due to the need for a full mesh topology within the chassis (every instance must be able to communicate with every other instance that is sharing the same interface). Therefore, there are limits to how many interfaces you can share.

In addition to the forwarding table, the chassis maintains a VLAN group table for VLAN subinterface forwarding. You can create up to 500 VLAN subinterfaces.

See the following limits for shared interface allocation:



Shared interface best practices

This topic outlines the best practices to follow when configuring shared interfaces to maximize performance and security.

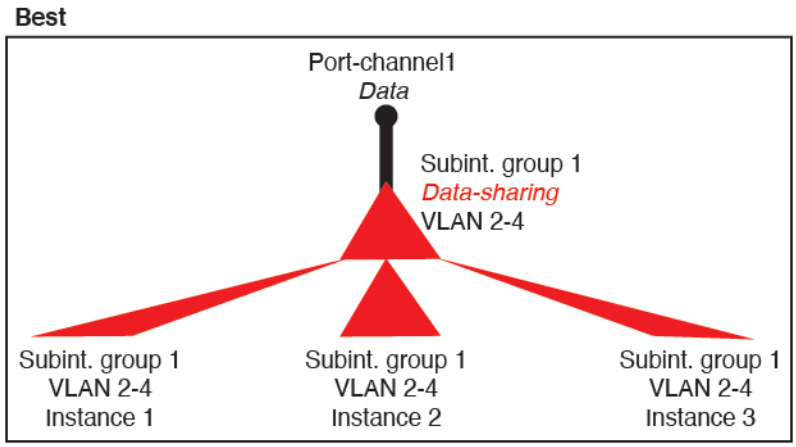
For optimal scalability of the forwarding table, share as few interfaces as possible. Instead, you can create up to 500 VLAN subinterfaces on one or more physical interfaces and then divide the VLANs among the container instances.

When sharing interfaces, follow these practices in the order of most scalable to least scalable:

1. Best: Share subinterfaces under a single parent, and use the same set of subinterfaces with the same group of instances.

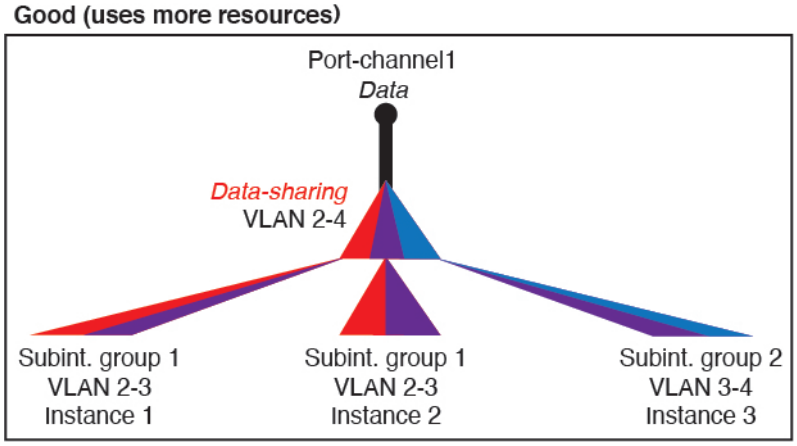
   For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel: Port-Channel1.2, 3, and 4 instead of Port-Channel2, Port-Channel3, and Port-Channel4. When you share subinterfaces from a single parent, the VLAN group table provides better scaling of the forwarding table than when sharing physical/EtherChannel interfaces or subinterfaces across parents.

Figure 39: Best: Shared Subinterface Group on One Parent



If you do not share the same set of subinterfaces with a group of instances, your configuration can cause more resource usage (more VLAN groups). For example, share Port-Channel1.2, 3, and 4 with instances 1, 2, and 3 (one VLAN group) instead of sharing Port-Channel1.2 and 3 with instances 1 and 2, while sharing Port-Channel1.3 and 4 with instance 3 (two VLAN groups).

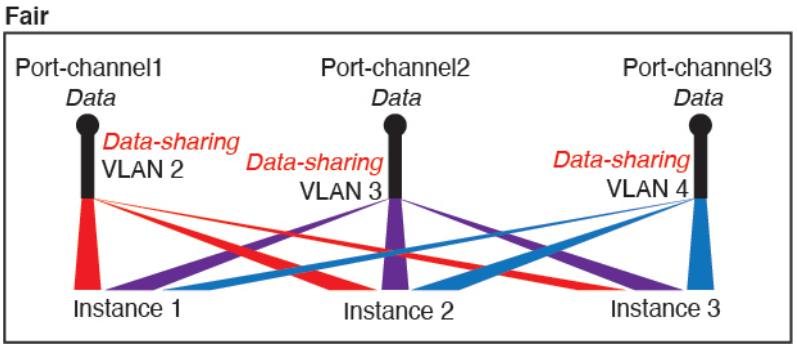
Figure 40: Good: Sharing Multiple Subinterface Groups on One Parent



2. Fair: Share subinterfaces across parents.

   For example, share Port-Channel1.2, Port-Channel2.3, and Port-Channel3.4 instead of Port-Channel2, Port-Channel4, and Port-Channel4. Although this usage is not as efficient as only sharing subinterfaces on the same parent, it still takes advantage of VLAN groups.

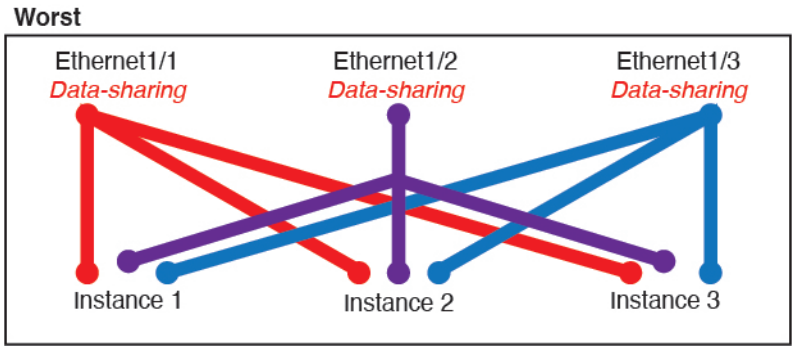
Figure 41: Fair: Shared Subinterfaces on Separate Parents



3. Worst: Share individual parent interfaces (physical or EtherChannel).

   This method uses the most forwarding table entries.

Figure 42: Worst: Shared Parent Interfaces



How the chassis classifies packets

The chassis processes incoming and outgoing packets as they pass through a device, ensuring that each packet is handled according to its type, destination, and security requirements.

Each packet that enters the chassis must be classified so that the chassis can determine to which instance to send a packet.

- Unique Interfaces: If only one instance is associated with the ingress interface, the chassis classifies the packet into that instance. For bridge group member interfaces (in transparent mode or routed mode), inline sets, or passive interfaces, this method is used to classify packets at all times.
- Unique MAC Addresses: The chassis automatically generates unique MAC addresses for all interfaces, including shared interfaces. If multiple instances share an interface, then the classifier uses unique MAC addresses assigned to the interface in each instance. An upstream router cannot route directly to an instance without unique MAC addresses. You can also set the MAC addresses manually when you configure each interface within the application.

Note If the destination MAC address is a multicast or broadcast MAC address, the packet is duplicated and delivered to each instance.

Classification examples

This reference provides examples to help you understand different classification scenarios and their implementation.

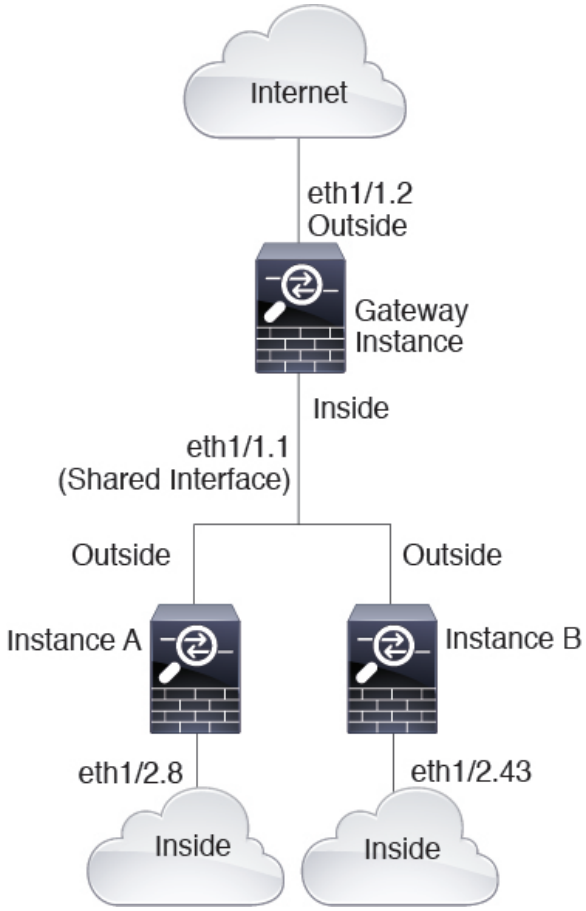
Cascading container instances

Cascading instances connect one instance's outside interface to another's inside interface, simplifying configuration by sharing parameters in the top instance. Avoid using shared interfaces with High Availability to prevent MAC address conflicts; use unique interfaces and an external switch instead.

Placing an instance directly in front of another instance is called *cascading instances*; the outside interface of one instance is the same interface as the inside interface of another instance. You might want to cascade instances if you want to simplify the configuration of some instances by configuring shared parameters in the top instance.

The following figure shows a gateway instance with two instances behind the gateway.

Figure 43: Cascading Instances



Note Do not use cascading instances (using a shared interface) with High Availability. After a failover occurs and the standby unit rejoins, MAC addresses can overlap temporarily and cause an outage. You should instead use unique interfaces for the gateway instance and inside instance using an external switch to pass traffic between the instances.

Typical multi-instance deployments

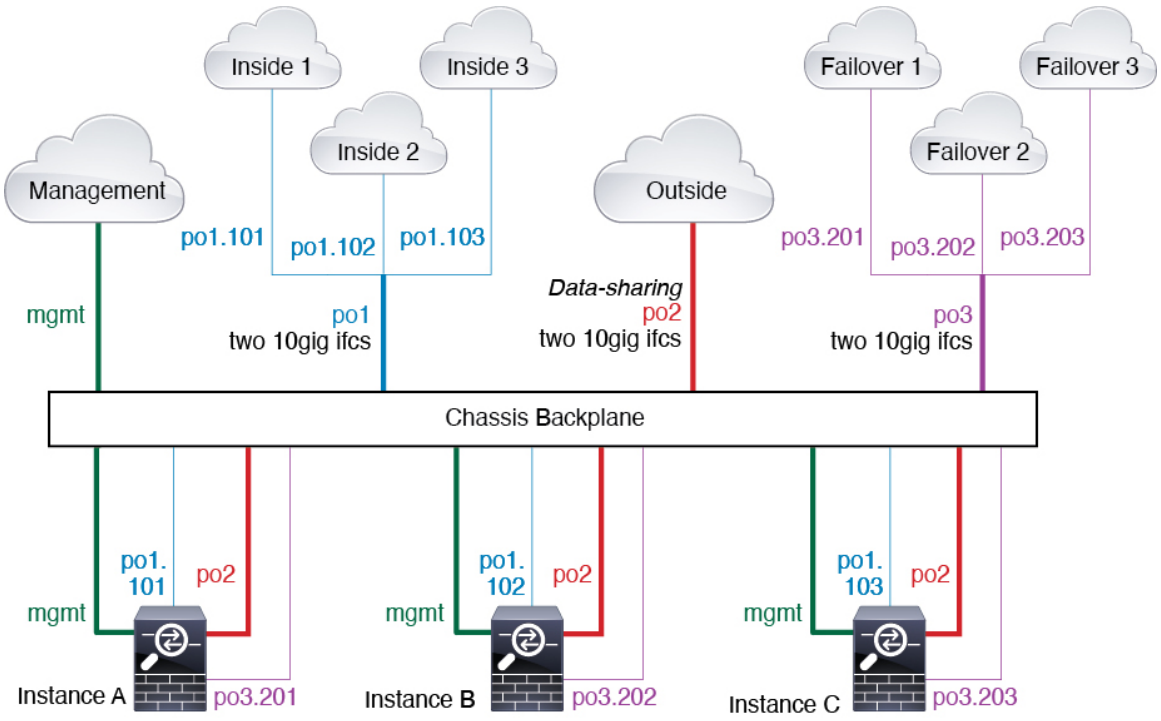
A typical multi-instance deployment is a firewall configuration that

- supports multiple container instances operating simultaneously,
- configures distinct interfaces for management, inside, outside, and failover connectivity, and
- utilizes dedicated subinterfaces and EtherChannels to enhance performance and reliability.

The following example includes three container instances in routed firewall mode. They include the following interfaces:

- Management: All instances and the chassis use the dedicated Management interface. Within each instance (and the chassis), the interface uses a unique IP address on the same management network.
- Inside: Each instance uses a subinterface on Port-Channel1 (data type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Each subinterface is on a separate network.
- Outside: All instances use the Port-Channel2 interface (data-sharing type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Within each application, the interface uses a unique IP address on the same outside network.
- Failover: Each instance uses a subinterface on Port-Channel3 (data type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Each subinterface is on a separate network.

Figure 44: Typical multi-instance deployment



Automatic MAC addresses for instance interfaces

An automatic MAC address for instance interfaces is a network identifier that is automatically generated and assigned to instance interfaces in multi-instance deployments.

The chassis automatically generates MAC addresses for instance interfaces, and guarantees that a shared interface in each instance uses a unique MAC address.

If you manually assign a MAC address to a shared interface within the instance, then the manually-assigned MAC address is used. If you later remove the manual MAC address, the autogenerated address is used. In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, we suggest that you manually set the MAC address for the interface within the instance.

Because autogenerated addresses start with A2, you should not start manual MAC addresses with A2 due to the risk of overlapping addresses.

The chassis generates the MAC address using the following format:

A2*xx.yyzz.zzzz*

Where *xx.yy* is a user-defined prefix or a system-defined prefix, and *zz.zzzz* is an internal counter generated by the chassis. The system-defined prefix matches the lower 2 bytes of the first MAC address in the burned-in MAC address pool that is programmed into the IDPROM. Use **connect fxos**, then **show module** to view the MAC address pool. For example, if the range of MAC addresses shown for module 1 is b0aa.772f.f0b0 to b0aa.772f.f0bf, then the system prefix will be f0b0.

The user-defined prefix is an integer that is converted into hexadecimal. For an example of how the user-defined prefix is used, if you set a prefix of 77, then the chassis converts 77 into the hexadecimal value 004D (*yyxx*). When used in the MAC address, the prefix is reversed (*xxyy*) to match the chassis native form:

A2**4D.00***zz.zzzz*

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03***zz.zzzz*

Performance scaling factors for multi-instance modes

A performance scaling factor is a throughput adjustment metric that

- enables calculation of effective throughput in multi-instance appliance modes,
- considers the percentage of CPU cores assigned to each instance, and
- accounts for efficiency reductions compared to single-instance (appliance mode) operation.

The maximum throughput (connections, VPN sessions) for a platform is calculated for an appliance mode device's use of memory and CPU (and this value is shown in **show resource usage**). If you use multiple instances, then you need to calculate the throughput based on the percentage of CPU cores that you assign to the instance. For example, if you use an instance with 50% of the cores, then you should initially calculate 50% of the throughput. Moreover, the throughput available to an instance may be less than that available to an appliance.

For detailed instructions on calculating the throughput for instances, refer to https://www.cisco.com/c/en/us/products/collateral/security/firewalls/white-paper-c11-744750.html.

Instances and high availability

A high availability pair is a deployment option that

- uses two instances on separate chassis to ensure service continuity,
- allows for multiple high availability pairs to be created from available instances, and
- permits standalone instances to exist alongside high availability pairs within the same chassis.

You can use High Availability using an instance on 2 separate chassis; for example, if you have 2 chassis, each with 10 instances, you can create 10 High Availability pairs. You can also have standalone instances on the same chassis as High Availability instances. For more information about the requirements, refer to Requirements and prerequisites for instances, on page 625.

Note Clustering is not supported.

Licenses for instances

This reference provides information about how licenses are assigned and consumed for chassis and feature instances.

All licenses are consumed per chassis and not per instance. Refer to these details.

- Essentials licenses are assigned to the chassis as a whole, one per chassis.
- Feature licenses are assigned to each instance, but you only *consume* one license per feature per chassis.

Requirements and prerequisites for instances

This reference provides the requirements and prerequisites for deploying instances on Secure Firewall models, including supported models, upgrade paths, resource limitations per model, software requirements, high availability requirements, and management center requirements.

Model support

- Secure Firewall 3110
- Secure Firewall 3120
- Secure Firewall 3130
- Secure Firewall 3140
- Secure Firewall 4215
- Secure Firewall 4225
- Secure Firewall 4245

Note The Secure Firewall 3105 is not supported.

Maximum container instances and resources per model

For each container instance, you can specify the number of CPU cores (or more specifically, threads) to assign to the instance. We use the term "core" generically to account for different hardware architecture. RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance.

Table 43: Maximum container instances and resources per model

Model	Max. Container Instances	Available CPU Cores (Threads)
Secure Firewall 3110	3	22
Secure Firewall 3120	5	30
Secure Firewall 3130	7	46
Secure Firewall 3140	10	62
Secure Firewall 4215	10	62
Secure Firewall 4225	21	126
Secure Firewall 4245	34	254

Software requirements

- You can run different versions of Firewall Threat Defense software on each instance as long as they are all compatible with the version of FXOS running on the chassis.
- You cannot deploy an instance with a patch version of the Firewall Threat Defense software because it is not a complete bundle. You need to first deploy the major or maintenance version, and then patch it after deployment.

High availability requirements

- The two instances in a High Availability configuration must:
  - Be on separate chassis.
  - Be on the same model.
  - Have the same interfaces assigned. All interfaces must be preconfigured in the chassis identically before you enable High Availability.
  - Use the same resource profile attributes. The profile names can be different, but the definitions need to match.

Cloud-Delivered Firewall Management Center requirements

For chassis management and all instances on the chassis, you must use the same Cloud-Delivered Firewall Management Center due to the licensing implementation.

Guidelines and limitations for instances

This reference provides guidelines and limitations for managing instances on chassis, including general requirements, interface configurations, VLAN subinterfaces, EtherChannels, data-sharing interfaces, and MAC address considerations.

General guidelines

- A single Cloud-Delivered Firewall Management Center must manage all instances on a chassis, as well as manage the chassis itself.
- For instances, these features are not supported:
  - TLS crypto acceleration
  - Clustering
  - Cloud-Delivered Firewall Management Center UCAPL/CC mode
  - Flow offload to hardware
  - LLDP
- Primary management of the chassis by Security Cloud Control cloud-delivered Cloud-Delivered Firewall Management Center and separate analytics-only management of the chassis by an on-prem Cloud-Delivered Firewall Management Center is not supported. You can however add Security Cloud Control-managed instances to an analytics-only on-prem Cloud-Delivered Firewall Management Center.

Management interface

- No support for a data interface for chassis management; only the dedicated Management interface can be used
- No DHCP addressing for the Management interface
- The management connection requires TCP port 8305; do not change the value from the default.

VLAN subinterfaces

- This document discusses *chassis* VLAN subinterfaces only. You can separately create subinterfaces within the instance.
- If you assign a parent interface to an instance, it only passes untagged (non-VLAN) traffic. Do not assign the parent interface unless you intend to pass untagged traffic.
- Subinterfaces are supported on Data or Data-sharing type interfaces.
- You can create up to 500 VLAN IDs.
- You cannot use subinterfaces for an inline set or as a passive interface.
- If you use a subinterface for the failover link, then all subinterfaces on that parent, and the parent itself, are restricted for use as failover links. You cannot use some subinterfaces as failover links, and some as regular data interfaces.

EtherChannels

- You can configure up to 48 EtherChannels, limited by the number of physical interfaces.
- The EtherChannel can have up to 8 active interfaces.
- All interfaces in the EtherChannel must be the same media type and speed capacity. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface

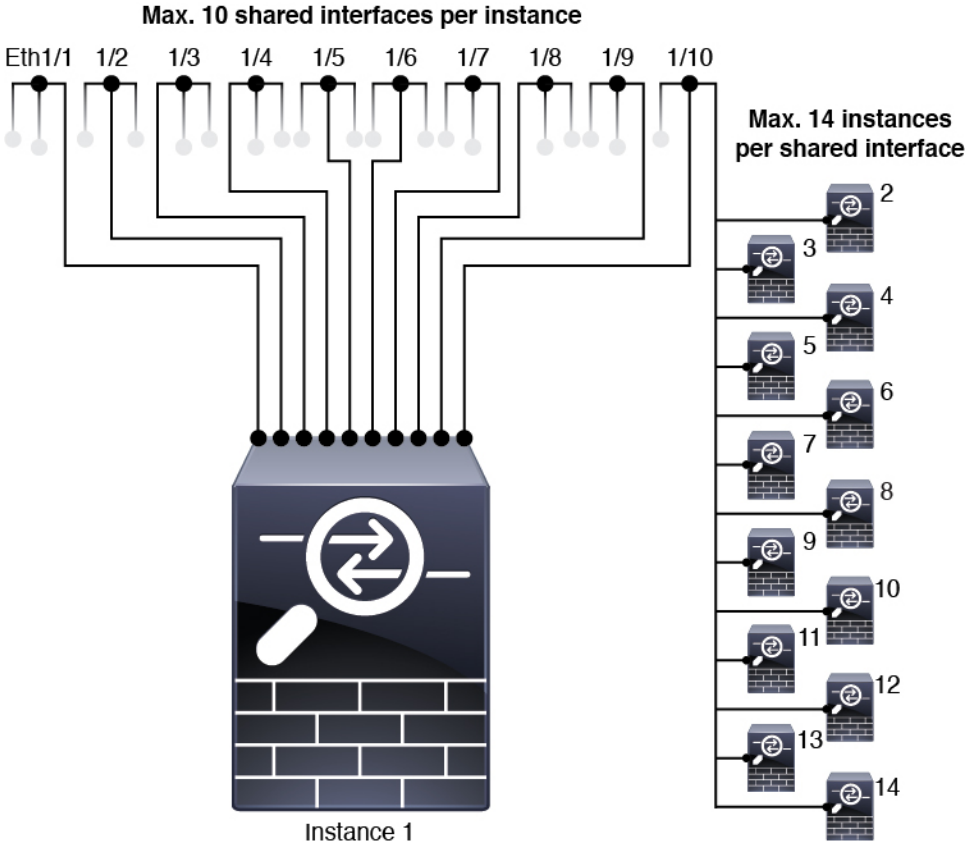
capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface, unless you set the speed to **Detect SFP**; in this case, you can use different interface capacities, and the lowest common speed is used.

- The chassis does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **VLAN dot1Q tag native** command, then the chassis will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch.
- In Cisco IOS software versions earlier than 15.1(1)S2, the chassis did not support connecting an EtherChannel to a switch stack. With default switch settings, if the chassis EtherChannel is connected cross stack, and if the primary switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-MAC persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.

Data-sharing interfaces

- Maximum 14 instances per shared interface. For example, you can allocate Ethernet1/1 to Instance1 through Instance14.

  Maximum 10 shared interfaces per instance. For example, you can allocate Ethernet1/1.1 through Ethernet1/1.10 to Instance1.



- You cannot use a data-sharing interface with a transparent firewall mode instance.
- You cannot use a data-sharing interface with inline sets or passive interfaces.

- You cannot use a data-sharing interface for the failover link.

Default MAC addresses

- MAC addresses for all interfaces are taken from a MAC address pool. For subinterfaces, if you decide to manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification. For more information, refer to Automatic MAC addresses for instance interfaces, on page 623.

Instance configuration

Prepare the chassis and management environment to enable instance configuration on your device.

Before you configure instances, you need to enable multi-instance mode, add the chassis to the Cloud-Delivered Firewall Management Center, and configure chassis interfaces. You can also customize the chassis settings.

Convert a device to Multi-Instance mode

Use this procedure to convert a device in the Cloud-Delivered Firewall Management Center to multi-instance mode. The conversion takes approximately 15 minutes. The system reboots and, as part of changing the mode, erases the configuration with the exception of the Management network settings and the admin password. The chassis hostname is set to "firepower-*model*." The Management IP address is assigned to the chassis for the management connection with the Cloud-Delivered Firewall Management Center. When you add instances, they will use separate IP addresses on the Management interface and maintain their own management connections.

After you convert to multi-instance mode, you can use the Cloud-Delivered Firewall Management Center to configure all chassis settings as well as instances. The Secure Firewall Chassis Manager or configuration at the FXOS CLI is not supported.

Note If you would prefer to use the CLI to convert to multi-instance mode before adding the device to the Cloud-Delivered Firewall Management Center, refer to Onboard a Chassis, on page 49.

Before you begin

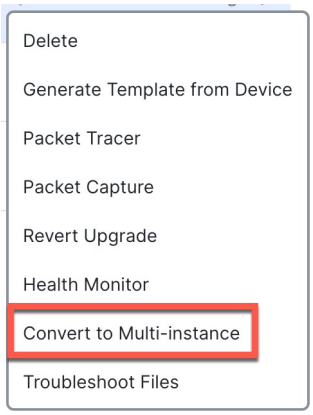
Add the appliance-mode device to the Cloud-Delivered Firewall Management Center as a standalone device. This device cannot use:

- A data interface for manager access
- DHCP for the Management interface
- Zero-Touch Provisioning

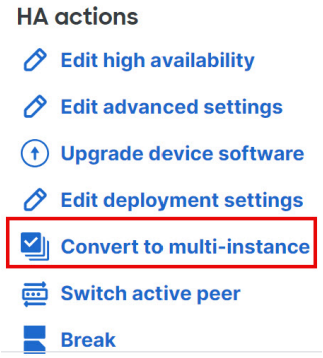
Procedure

Step 1 On **Devices** > **Device Management**, next to the device you want to convert to multi-instance mode, choose **More** (⋮) > **Convert to Multi-Instance**.

Figure 45: Convert to Multi-Instance

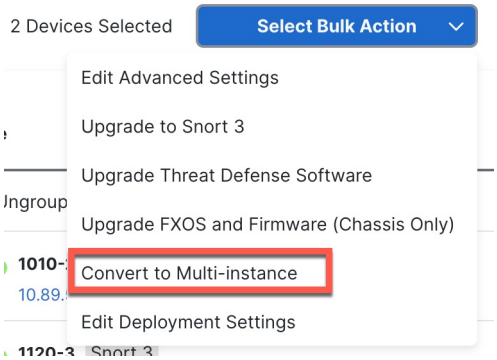


If you have switched to the new device management page, select your device to open the right-pane. Under **HA actions**, select **Convert to multi-instance**.



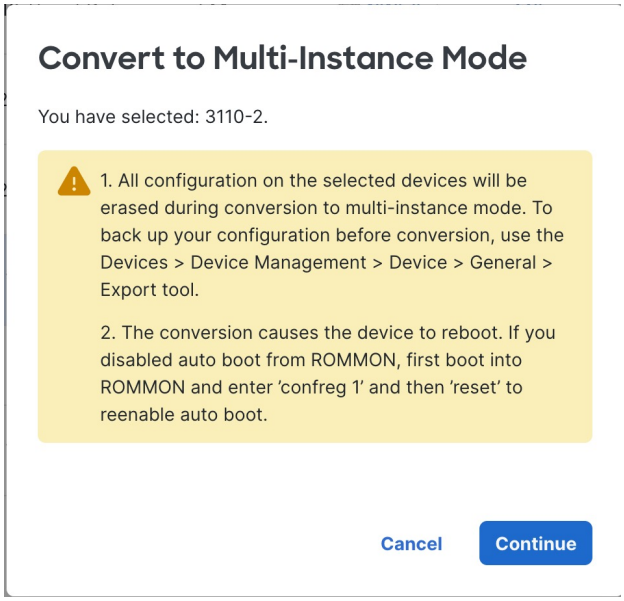
You can alternatively check the check-box next to multiple devices that you want to convert, and choose **Select Bulk Action** > **Convert to Multi-Instance**.

Figure 46: Bulk Conversion



Step 2 Confirm that you want to perform the conversion and click **Continue**.

Figure 47: Conversion Confirmation

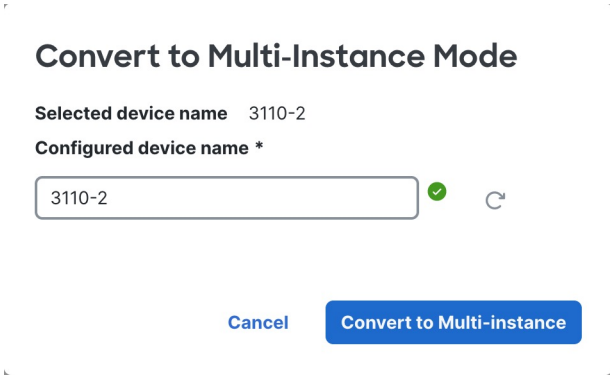


A readiness check is performed. The check might fail if, for example, a deployment is in progress.

Step 3 Optionally change the name of the chassis, and click **Convert to Multi-Instance**. By default, the device name is used, but you may want a different naming convention for multi-instance chassis.

Note
If the readiness check failed, for example because of a deployment in progress, you can wait for the process to complete and click ↻ to rerun the readiness check so you can continue.

Figure 48: Rename Chassis

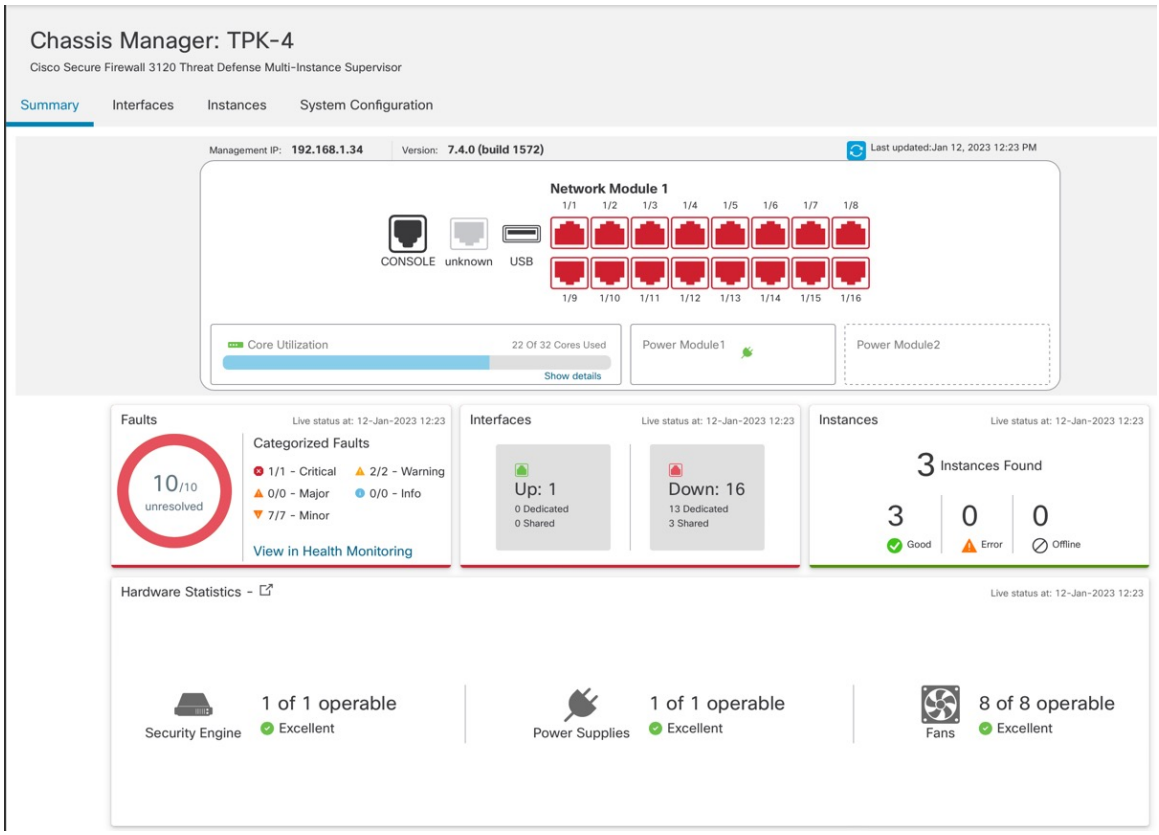


Wait for approximately 15 minutes, during which the device is removed from the device list, and then after conversion, it is re-added as a chassis.

Step 4 To view and configure the chassis, click **Manage** in the **Chassis** column, or click **Edit** ( ).

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Figure 49: Chassis Summary



Chassis interface configuration

A chassis interface is a physical network interface that

- enables connection to the firewall's internal and external networks,
- supports configuration of basic Ethernet settings, VLAN subinterfaces, and EtherChannel grouping, and
- defaults to a disabled state until explicitly enabled by an administrator.

Additional information

Note To configure breakout ports and perform other network module operations, refer to Network module management for Secure Firewall 3100/4200, on page 1015.

Note For information about the **Sync Device** button, refer to Sync Interface Changes with the Cloud-Delivered Firewall Management Center, on page 1014.

Configure a physical interface

You can physically enable and disable interfaces, as well as set the interface speed and duplex and other hardware settings. To use an interface, it must be physically enabled for the chassis and logically enabled in the instance. By default, physical interfaces are disabled. For VLAN subinterfaces, the admin state is inherited from the parent interface.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

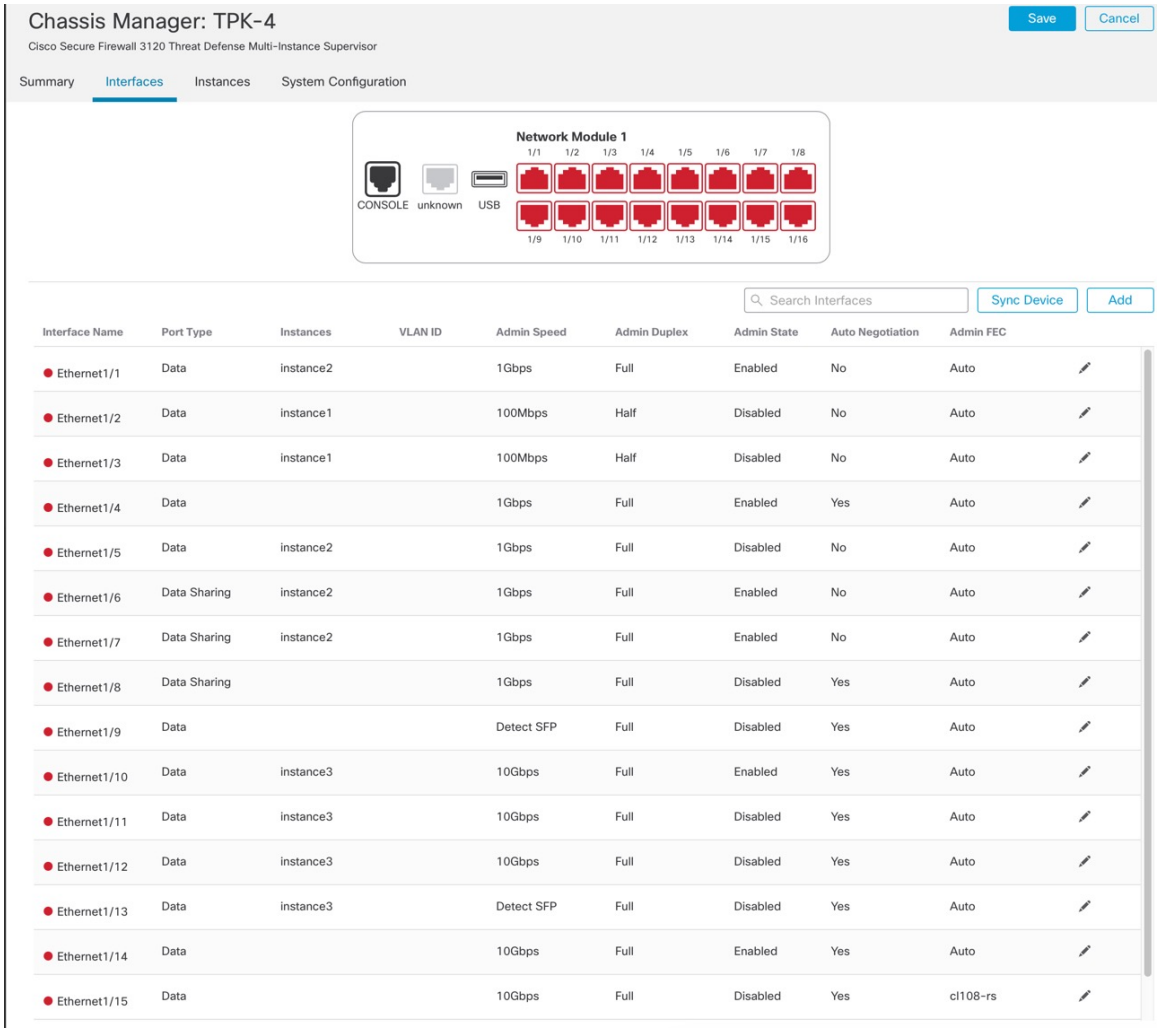
Figure 50: Manage Chassis



The **Chassis Manager** page opens for the chassis to the **Summary** page.

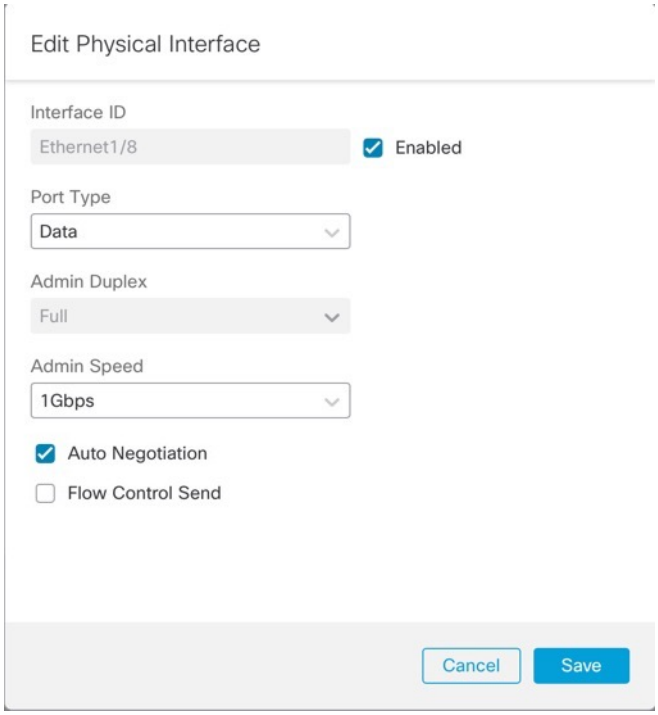
Step 2 Click **Interfaces**.

Figure 51: Interfaces



Step 3 Click **Edit** (✎) for the interface you want to edit.

Figure 52: Edit Physical Interface



Step 4 Enable the interface by checking the **Enabled** check box.

- For the **Port Type**, choose **Data** or **Data Sharing**.

  Figure 53: Port Type

- Set the **Admin Duplex**.

  Speeds of 1Gbps and higher only support **Full** duplex. SFP interfaces only support **Full** duplex.

- Set the **Admin Speed**.

  For SFPs, Choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always Full, and auto-negotiation is always enabled. This option is useful if you later change the network module to a different model, and want the speed to update automatically.

Step 5 (Optional) Check **Flow Control Send** to enable pause (XOFF) frames for flow control.

Flow control enables connected Ethernet ports to control traffic rates during congestion by allowing congested nodes to pause link operation at the other end. If the threat defense port experiences congestion (exhaustion of queuing resources on the internal switch) and cannot receive any more traffic, it notifies the other port by sending a pause frame to stop sending until the condition clears. Upon receipt of a pause frame, the sending device stops sending any data packets, which prevents any loss of data packets during the congestion period.

Note
The Firewall Threat Defense supports transmitting pause frames so that the remote peer can rate-control the traffic.

However, receiving of pause frames is not supported.

The internal switch has a global pool of 8000 buffers of 250 bytes each, and the switch allocates buffers dynamically to each port. A pause frame is sent out every interface with flow control enabled when the buffer usage exceeds the global high-water mark (2 MB (8000 buffers)); and a pause frame is sent out of a particular interface when its buffer exceeds the port high-water mark (.3125 MB (1250 buffers)). After a pause is sent, an XON frame can be sent when the buffer usage is reduced below the low-water mark (1.25 MB globally (5000 buffers); .25 MB per port (1000 buffers)). The link partner can resume traffic after receiving an XON frame.

Only flow control frames defined in 802.3x are supported. Priority-based flow control is not supported.

Step 6 (Optional) Check **Auto Negotiation** to set the interface to negotiate the speed, link status, and flow control. You cannot edit this setting for speeds lower than 1Gbps. For SFP interfaces, you can only disable auto-negotiation when the speed is set to 1Gbps.

Step 7 Click **Save** and then **Save** in the top right of the **Interfaces** page.

You can now **Deploy** the policy to the chassis. The changes are not active until you deploy them.

Configure an EtherChannel

An EtherChannel (also known as a port channel) can include up to 8 member interfaces of the same media type and capacity, and must be set to the same speed and duplex. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface, unless you set the speed to **Detect SFP**; in this case, you can use different interface capacities, and the lowest common speed is used.

The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices. LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. "On" mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

When the chassis creates an EtherChannel, the EtherChannel stays in a **Suspended** state for Active LACP mode or a **Down** state for On LACP mode until you assign it to a logical device, even if the physical link is up. The EtherChannel will be brought out of this **Suspended** state when it is added to an instance.

Before you begin

Enable physical interfaces and set hardware parameters. For more information, refer to Configure a physical interface, on page 633.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

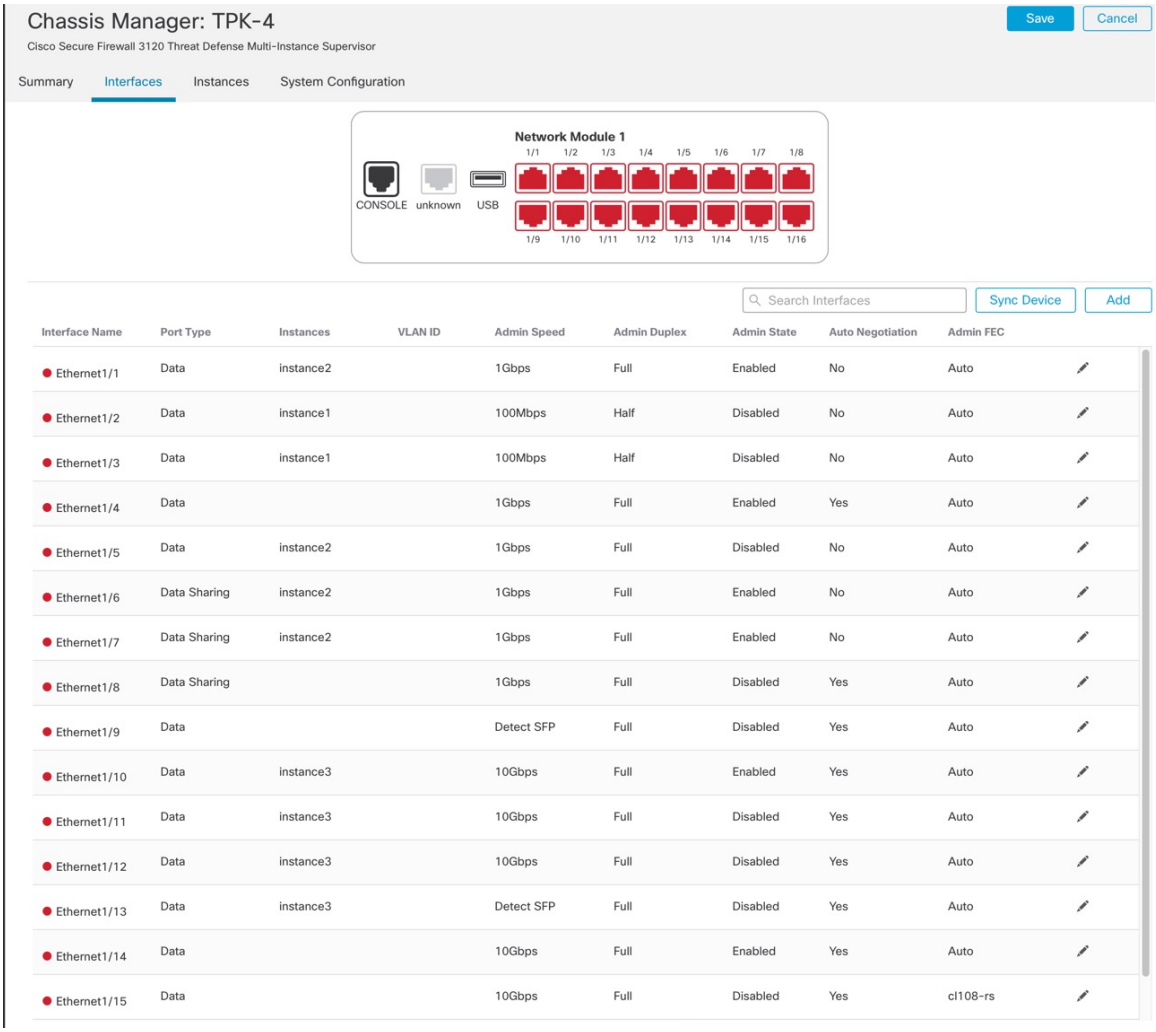
Figure 54: Manage Chassis



The **Chassis Manager** page opens for the chassis to the **Summary** page.

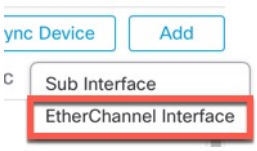
Step 2 Click **Interfaces**.

Figure 55: Interfaces



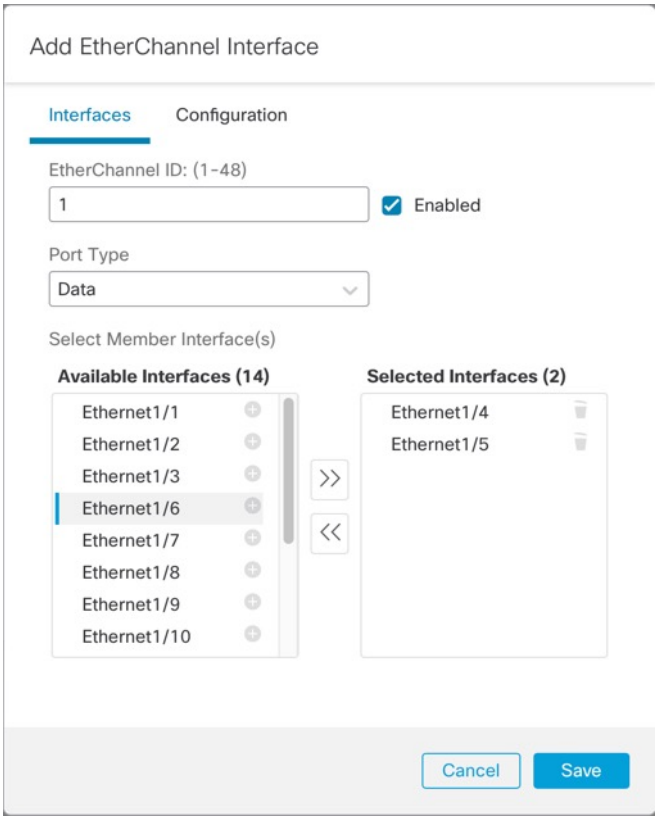
Step 3 Click **Add** > **EtherChannel Interface**.

Figure 56: Add EtherChannel



Step 4 Set the following **Interfaces** parameters.

Figure 57: Interfaces settings



a) For the **EtherChannel ID**, specify an ID between 1 and 48.
b) Check **Enabled**.
c) For the **Port Type**, choose **Data** or **Data Shared**.

   For information about the port type, refer to Interface types, on page 615.

d) To add a physical interface to the EtherChannel, select click **Add** (⊕) in the **Available Interfaces** list to move it to the **Selected Interfaces** list.

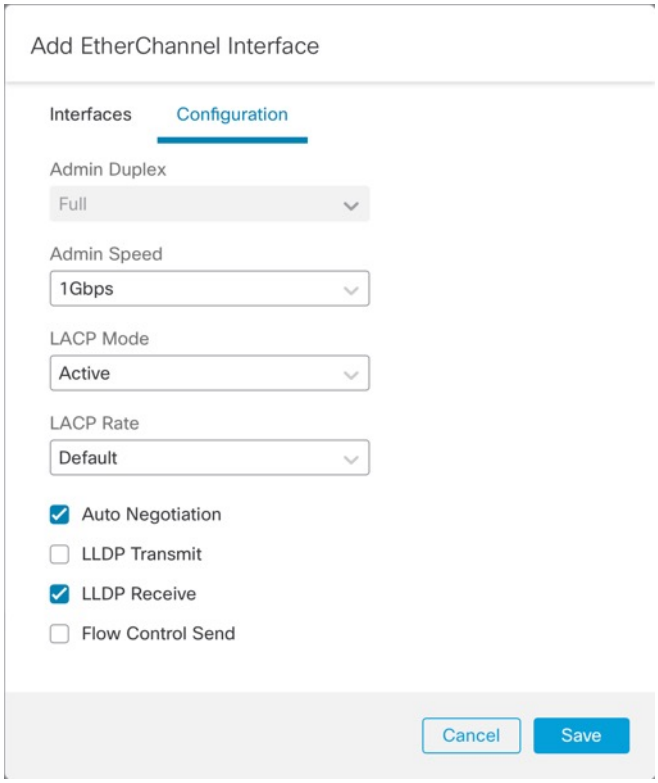
   To add or remove all interfaces, click the double arrow button.

   Note
   You cannot add an interface that is already assigned to an instance.

Step 5 (Optional) Set the following **Configuration** parameters.

Many of these settings (excluding the LACP settings) set the requirements for interfaces to be included in the EtherChannel; they do not override the settings of member interfaces. So if you check **LLDP Transmit**, for example, you should only add interfaces that have that setting. If you set the **Admin Speed** to 1 Gbps, then only 1Gbps interfaces can be included.

Figure 58: Configuration settings



a) Choose the required **Admin Duplex** for the member interfaces, **Full Duplex** or **Half Duplex**.

If you add a member interface that is configured with the specified duplex, it will not successfully join the port channel.

b) Choose the required **Admin Speed** for the member interfaces from the drop-down list.

If you add a member interface that is not at the specified speed, it will not successfully join the port channel.

c) Choose the **LACP Mode**, **Active** or **On**.

- **Active**—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- **On**—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

Note
It may take up to three minutes for an EtherChannel to come up to an operational state if you change its mode from On to Active or from Active to On.

d) Choose the **LACP Rate**, **Default**, **Fast**, or **Normal**.

The default is **Fast**.

e) Choose the required Link Layer Discovery Protocol (LLDP) settings for member interfaces by checking **LLDP Transmit** and/or **LLDP Receive**.
f) Check the required **Flow Control Send** setting for member interfaces.

Step 6 Click **Save** and then **Save** in the top right of the **Interfaces** page.

You can now **Deploy** the policy to the chassis. The changes are not active until you deploy them.

Configure a subinterface

You can add up to 500 subinterfaces to your chassis.

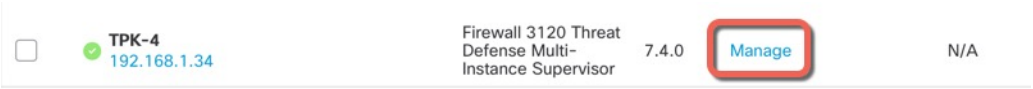
VLAN IDs per interface must be unique, and within an instance, VLAN IDs must be unique across all assigned interfaces. You can reuse VLAN IDs on *separate* interfaces as long as they are assigned to different instances. However, each subinterface still counts towards the limit even though it uses the same ID.

This section discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the instance. For more information, refer to Chassis interfaces vs. instance interfaces, on page 615.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

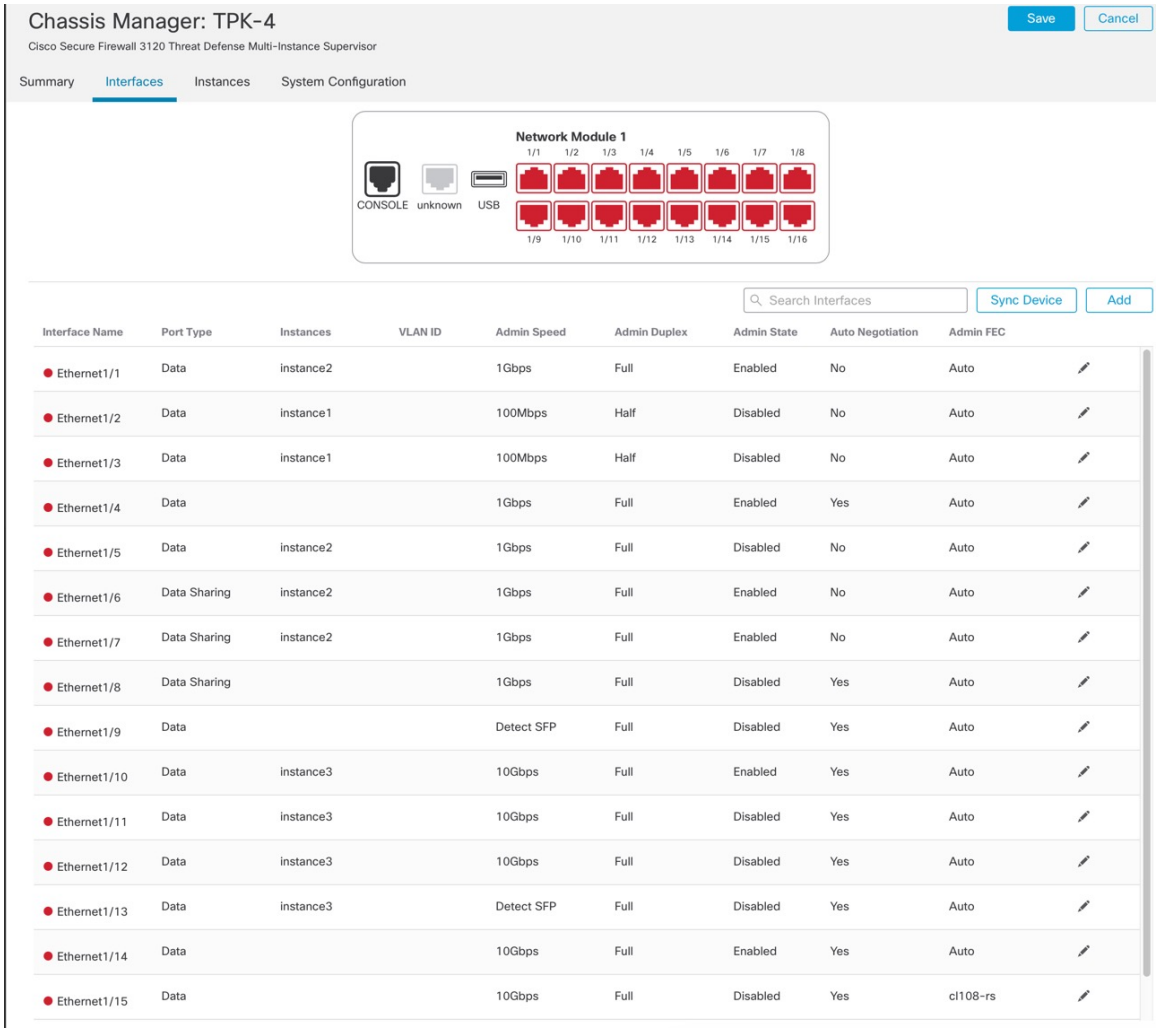
Figure 59: Manage Chassis



The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **Interfaces**.

Figure 60: Interfaces



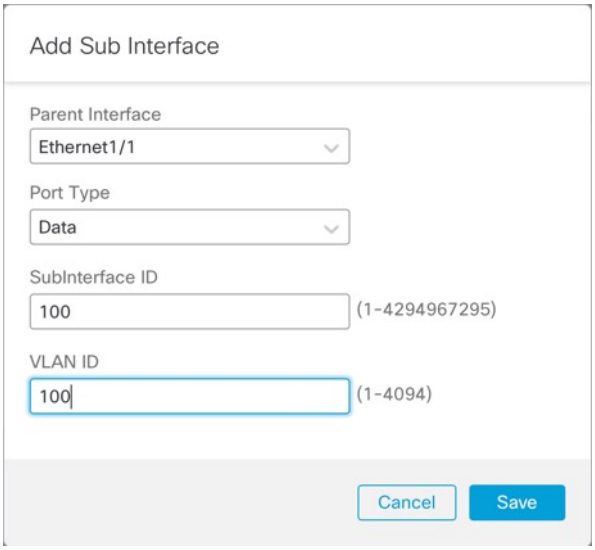
Step 3 Click **Add** > **Subinterface**.

Figure 61: Add subinterface



Step 4 Set the following parameters.

Figure 62: Subinterface settings



Step 5 Click **Save** and then **Save** in the top right of the **Interfaces** page.

You can now **Deploy** the policy to the chassis. The changes are not active until you deploy them.

Add an instance

Use this task when you need to add one or more instances to a chassis already in multi-instance mode. The number of supported instances depends on your model; refer to Requirements and prerequisites for instances, on page 625.

Before you begin

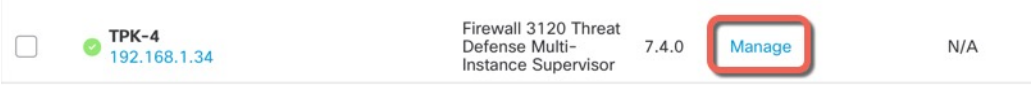
Onboard the multi-instance chassis using the CLI, on page 667.

Follow these steps to add an instance to a chassis in multi-instance mode:

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

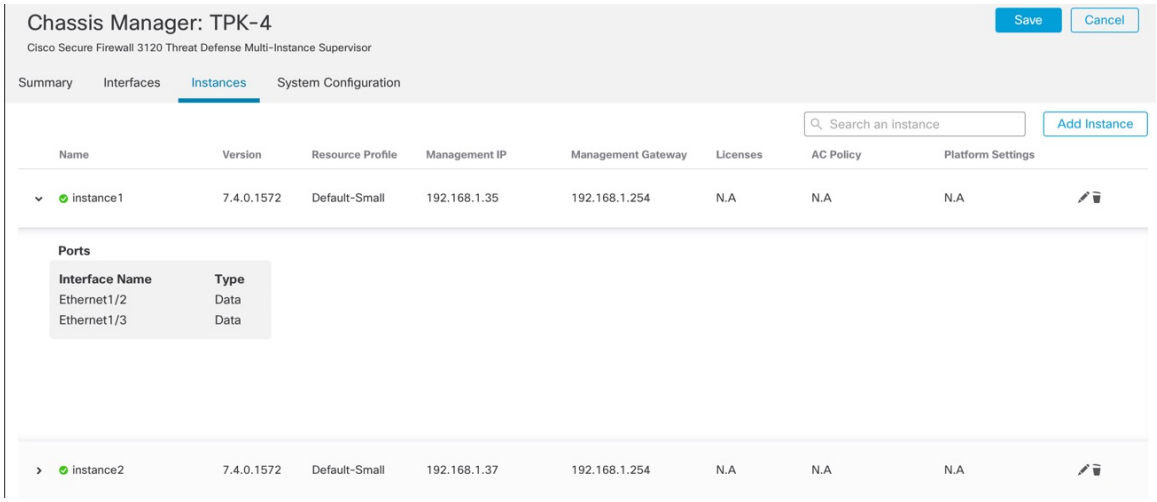
Figure 63: Manage Chassis



The **Chassis Manager** page opens for the chassis to the **Summary** page.

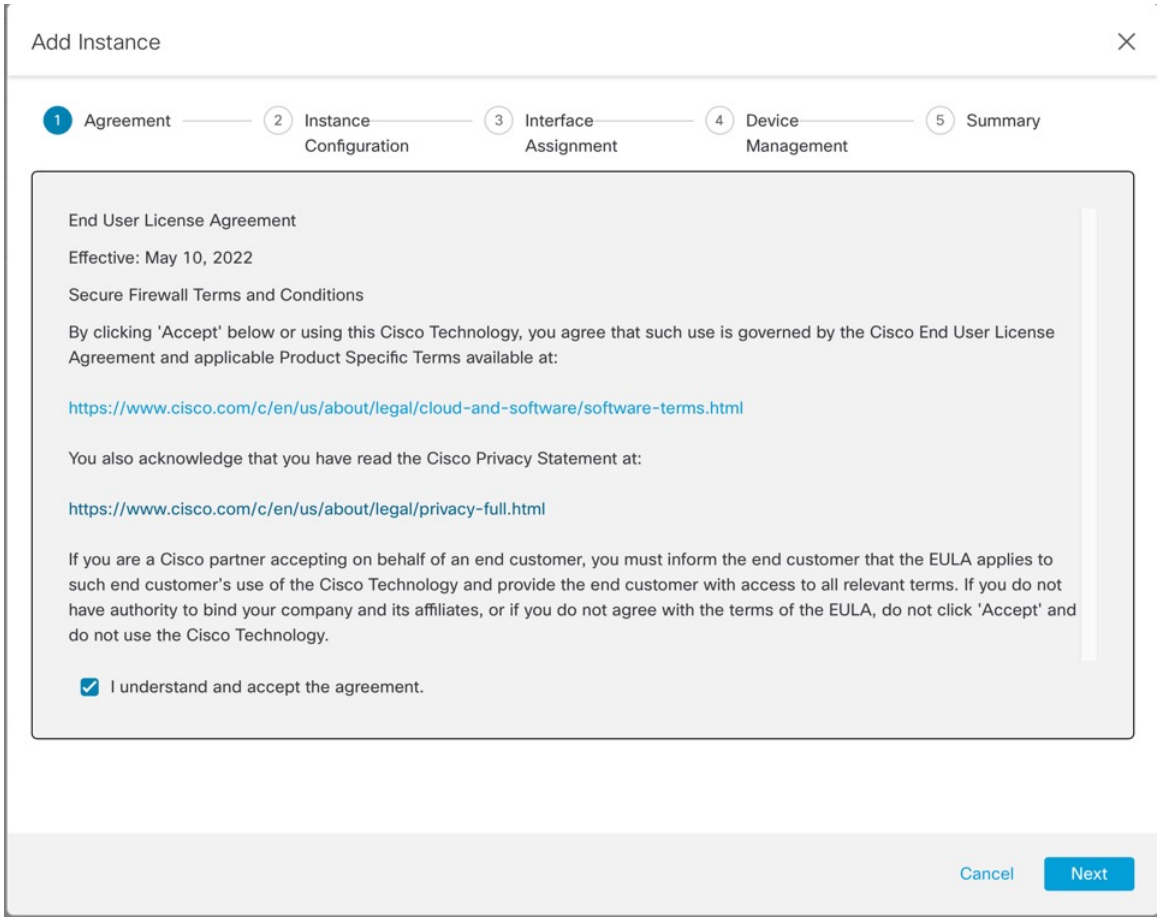
Step 2 Click **Instances**, and click **Add Instance**.

Figure 64: Instances



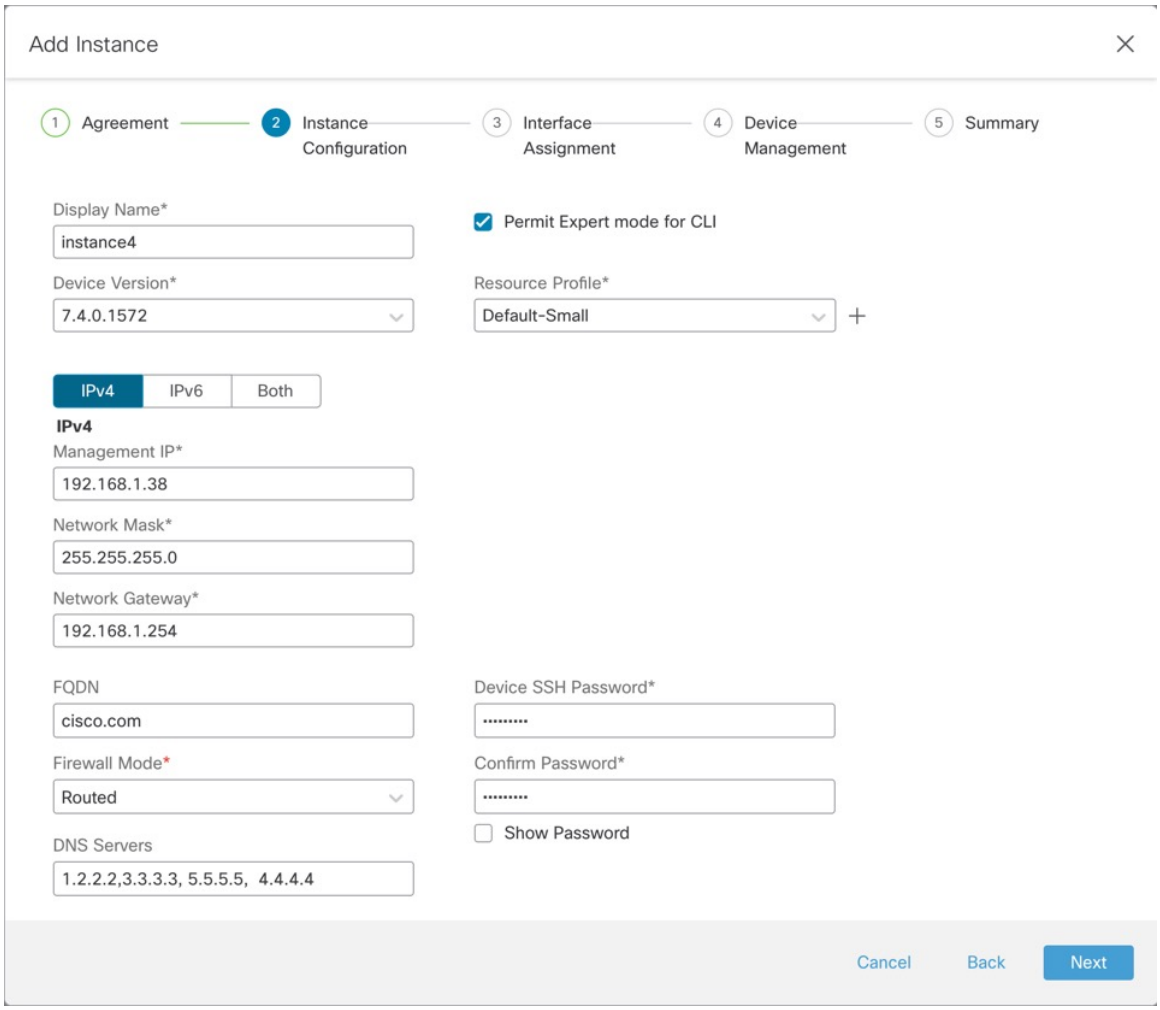
Step 3 On **Agreement**, check **I understand and accept the agreement**, then click **Next**.

Figure 65: Agreement



Step 4 On **Instance Configuration**, set the instance parameters, then click **Next**.

Figure 66: Instance configuration



- **Display Name**
- **Device Version**: Versions listed are packages currently downloaded to the chassis. Patch versions are not listed and cannot be used because they don't contain the entire bundle. To upgrade to a new package, refer to **Devices** > **Upgrade** > **Chassis Upgrade**. When you upgrade, both the old Firewall Threat Defense version and the new Firewall Threat Defense version will be listed in the menu. To download an older package, you need to use the FXOS CLI. Note: Both FXOS and Firewall Threat Defense images are included in the same package. For more information, refer to troubleshooting guide.

  For example:

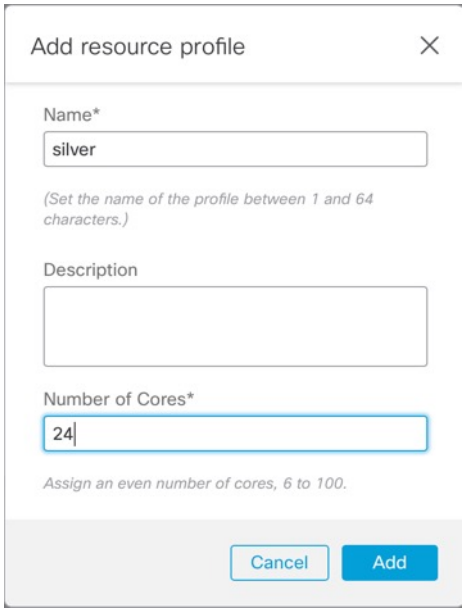
```
firepower-3110# scope firmware
firepower-3110# download image
https://10.10.7.89/Cisco_FTD_SSP_FP3K_Upgrade-7.4.1-1.sh.DEV.tar
Please use the command 'show download-task' or 'show download-task detail' to check
download progress.
% Download-task Cisco_FTD_SSP_FP3K_Upgrade-7.4.1-1.sh.DEV.tar : completed successfully.
```

- **IPv4**, **IPv6**, or **Both**: Set a **Management IP** address on the same network as the chassis Management interface. Set the **Network Mask** and gateway (likely the same gateway as the chassis). The chassis Management interface is shared with each instance, and each instance has its own IP address on the network. You can SSH to this IP address by default to reach the Firewall Threat Defense CLI.
- (Optional) **FQDN**
- **Firewall Mode**: **Routed** or **Transparent**. For more information about the firewall mode, refer to Transparent or Routed Firewall Mode, on page 153.
- **DNS Servers**: Enter a comma-separated list of DNS servers for management traffic only.
- (Optional) **Permit Expert Mode for CLI**: Expert Mode provides Firewall Threat Defense shell access for advanced troubleshooting.

  If you enable this option, then users who access the instance directly from an SSH session can enter Expert Mode. If you disable this option, then only users who access the instance from the FXOS CLI can enter Expert Mode. We recommend disabling this option to increase isolation between instances.

  Use Expert Mode only if a documented procedure tells you it is required, or if the Cisco Technical Assistance Center asks you to use it. To enter this mode, use the **expert** command in the Firewall Threat Defense CLI.
- **Resource Profile**: The resource profile sets the number of CPU cores; RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance. The chassis includes these default resource profiles: Default-Small, Default-Medium, and Default-Large. You can add additional profiles for this chassis by clicking **Add** (+). You cannot later edit the resource profile.

Figure 67: Add resource profile



  - The minimum number of cores is 6.

  Note

Instances with a smaller number of cores might experience relatively higher CPU utilization than those with larger numbers of cores. Instances with a smaller number of cores are more sensitive to traffic load changes. If you experience traffic drops, try assigning more cores.

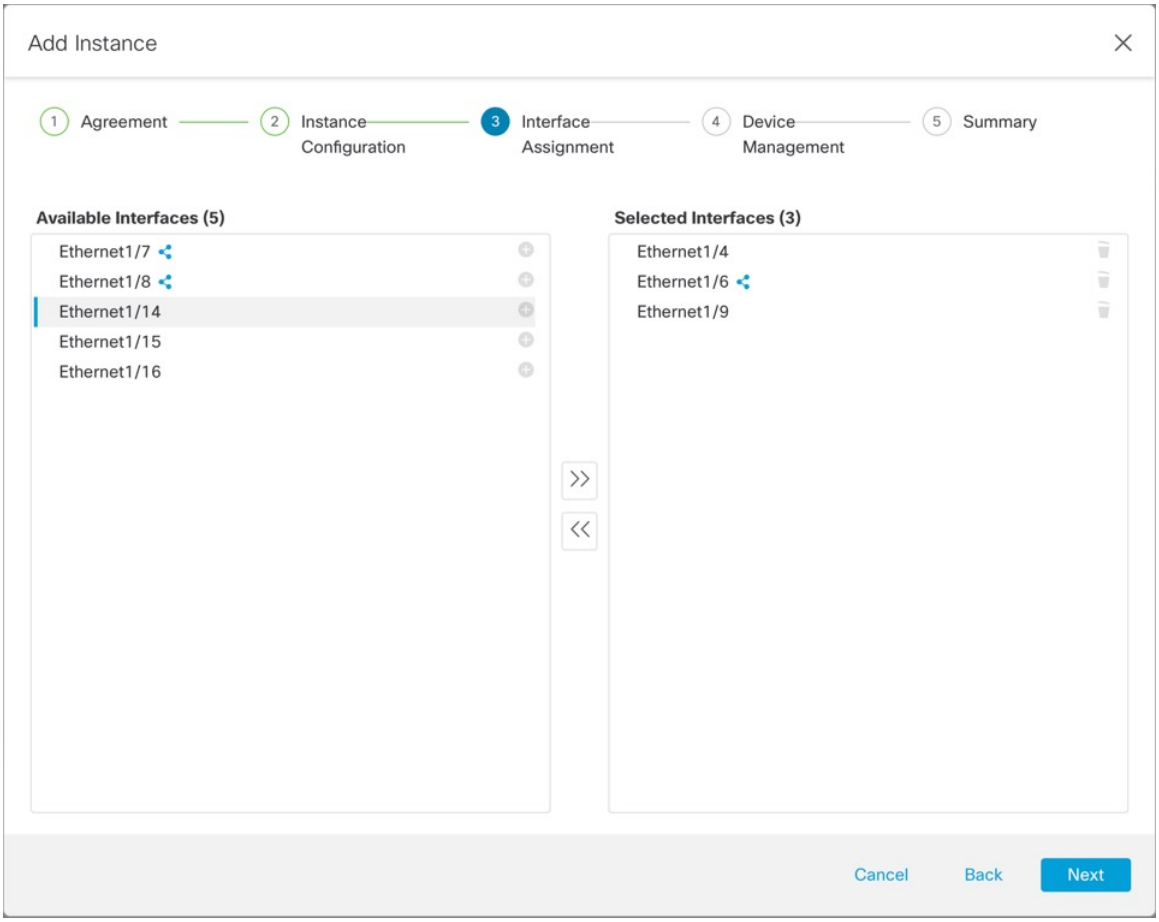
- You can assign cores as an even number (6, 8, 10, 12, 14 etc.) up to the maximum.
- The maximum number of cores available depends on the model; refer to Requirements and prerequisites for instances, on page 625.

If you later assign a different resource profile, then the instance will reload, which can take approximately 5 minutes. Note that for an established High Availability pair , if you assign a different-sized resource profile, be sure to make all members the same size as soon as possible.

- **Device SSH Password**: Set the Firewall Threat Defense admin user password for CLI access, either SSH or console. Repeat the password in the **Confirm Password** field.

Step 5 On **Interface Assignment**, assign the chassis interfaces to the instance, then click **Next**.

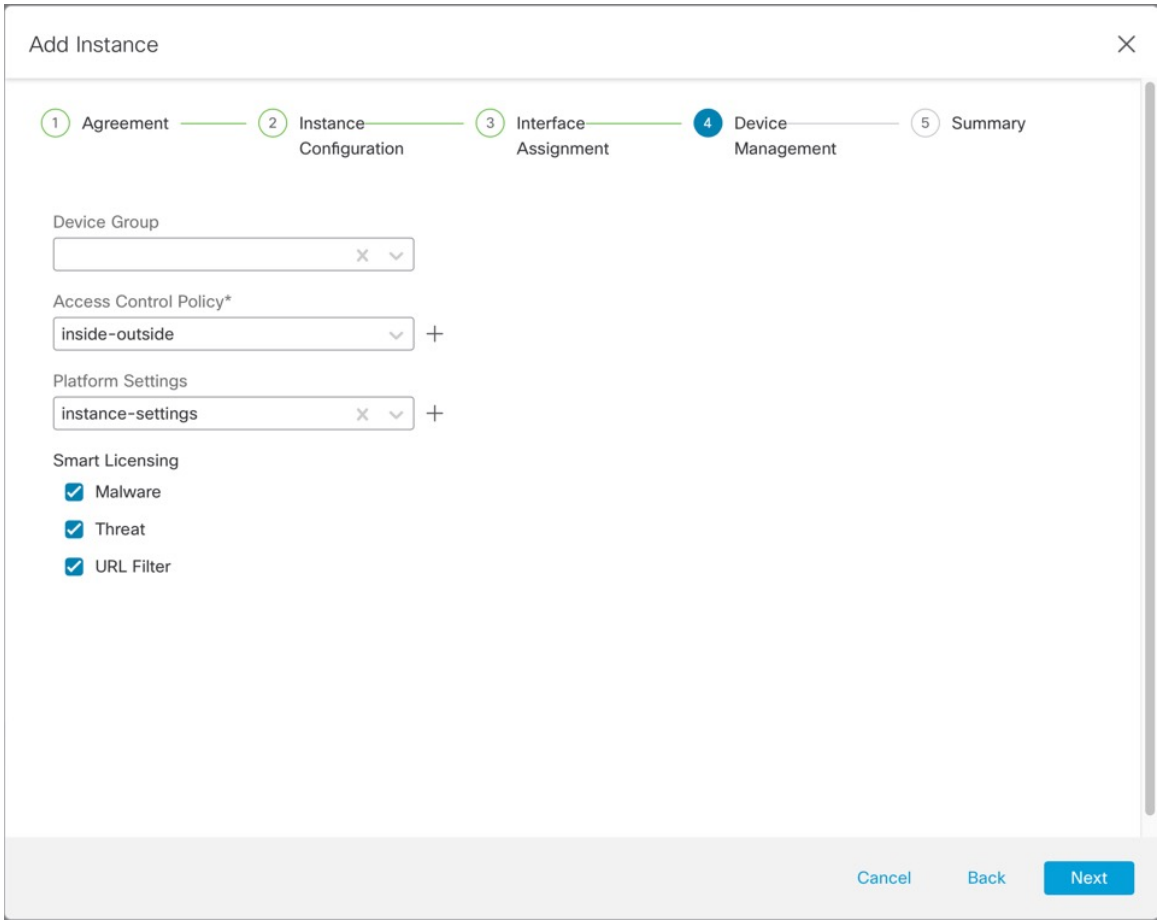
Figure 68: Interface assignment



Shared interfaces show the sharing icon ( ).

Step 6 On **Device Management**, set the device-specific settings, then click **Next**.

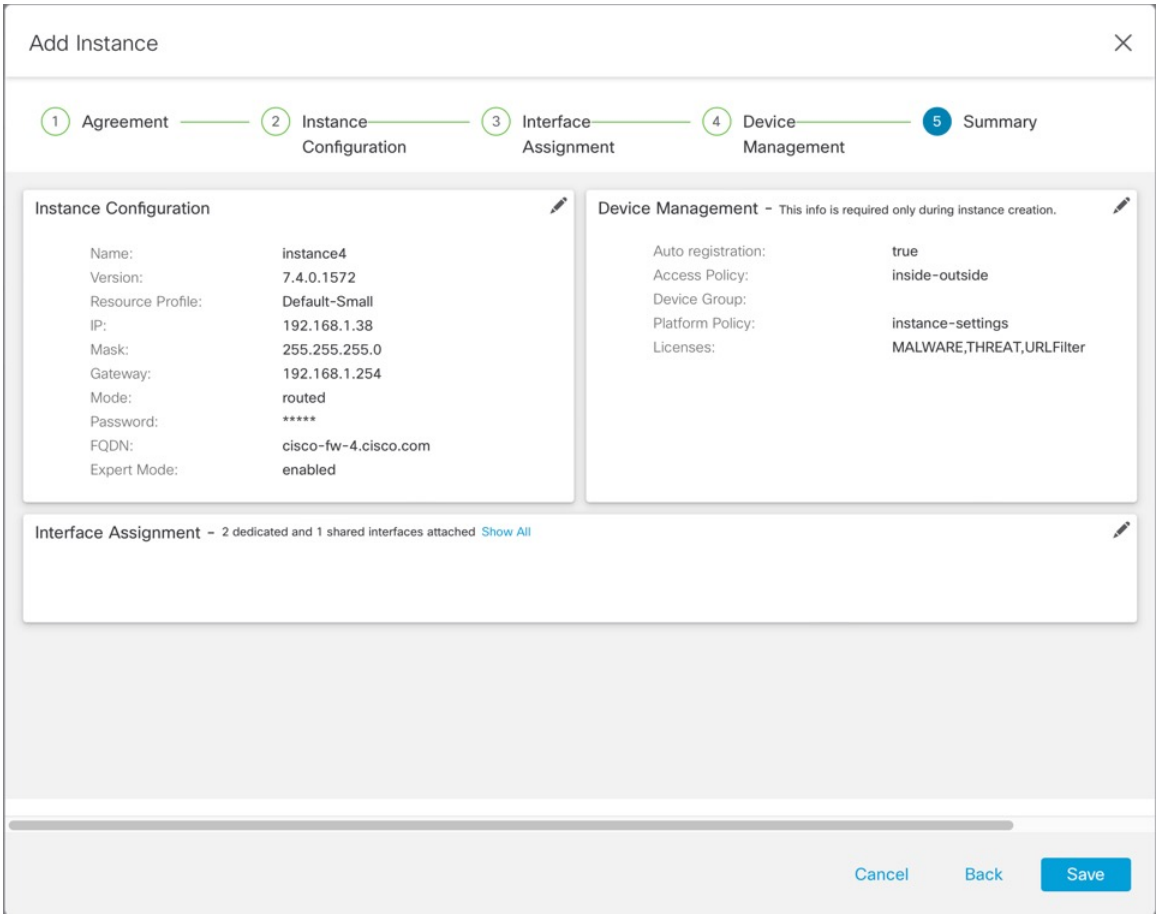
Figure 69: Device management



- **Device Group**
- **Access Control Policy**: Choose an existing access control policy, or create a new policy.
- **Platform Settings**: Choose an existing platform setting policy, or create a new policy.
- **Smart Licensing**

Step 7 On **Summary**, confirm your settings, then click **Save**.

Figure 70: Summary



You can edit any settings on this screen before saving the instance. After you save, the instance is added to the **Instances** screen.

Step 8 On the **Instances** screen, click **Save**.

Deploy the chassis configuration.

Step 9 Manually add the instance to Security Cloud Control as an FTD.

a) Add the instance **Using a CLI Registration Key**. For more information, refer to Onboard a Device with a CLI Registration Key, on page 22.

For the **Device Name**, you can use the instance name that you specified or any other name.

b) Access the chassis CLI and connect to the instance CLI. For more information, refer to Log into the Command-Line Interface on the device, on page 99.

connect FTD *name*

c) Delete the current manager; this manager needs to be deleted and reconfigured.

configure manager delete

d) Paste the generated **configure manager add** command from Security Cloud Control.

System configuration customization

You can configure chassis-level settings such as SNMP. You can also import or export the chassis FXOS configuration.

Configure SNMP

You can access chassis-level MIBs through the data interface of one of the instances, which you specify in the chassis system configuration. You can only use this instance for chassis SNMP information. You cannot access SNMP through the chassis Management interface.

Before you begin

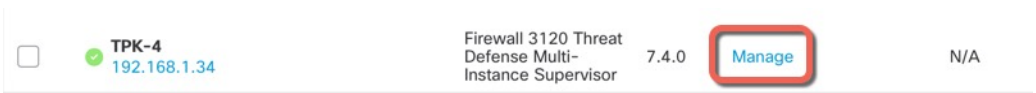
Configure SNMP for one of the instances. For more information, refer to Configure SNMP for a security device, on page 1264.

Follow these steps to configure SNMP:

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

Figure 71: Manage Chassis

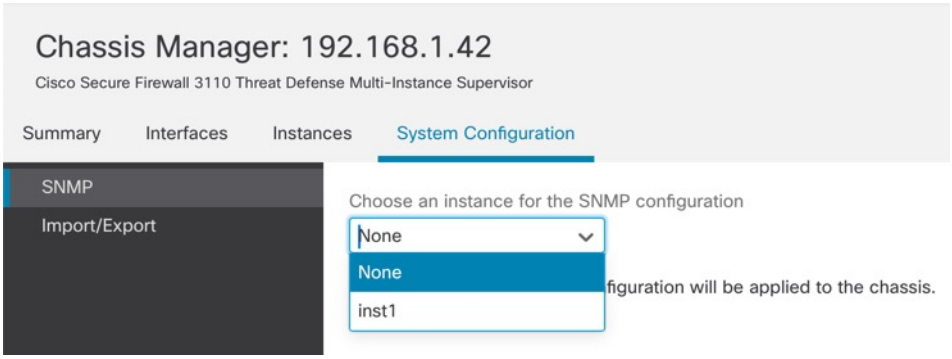


The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **System Configuration**.

Step 3 Click **SNMP**, and choose the instance from the drop-down list.

Figure 72: SNMP



SNMP on the chassis will be accessible from the selected instance.

Step 4 Click **Save**.

Step 5 Deploy the chassis configuration.

Import or export the chassis configuration

You can use the configuration export feature to export an XML file containing chassis configuration settings to your local computer. You can later import that configuration file to quickly apply the configuration settings to your chassis to return to a known good configuration or to recover from a system failure.

Note If you need to import the configuration into a new chassis of the same model, contact Cisco TAC.

When exporting, only chassis configurations are exported; instance configuration settings are not exported. Instances need to be backed up separately using the device backup/restore feature.

When importing, all existing configurations on the chassis will be replaced by the configuration in the import file.

Before you begin

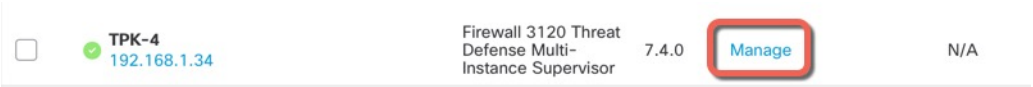
When you import the configuration to the same chassis, make sure the chassis matches for the following:

- Same chassis software version
- Same Firewall Threat Defense instance images
- Same network modules

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

Figure 73: Manage Chassis



The **Chassis Manager** page opens for the chassis to the **Summary** page.

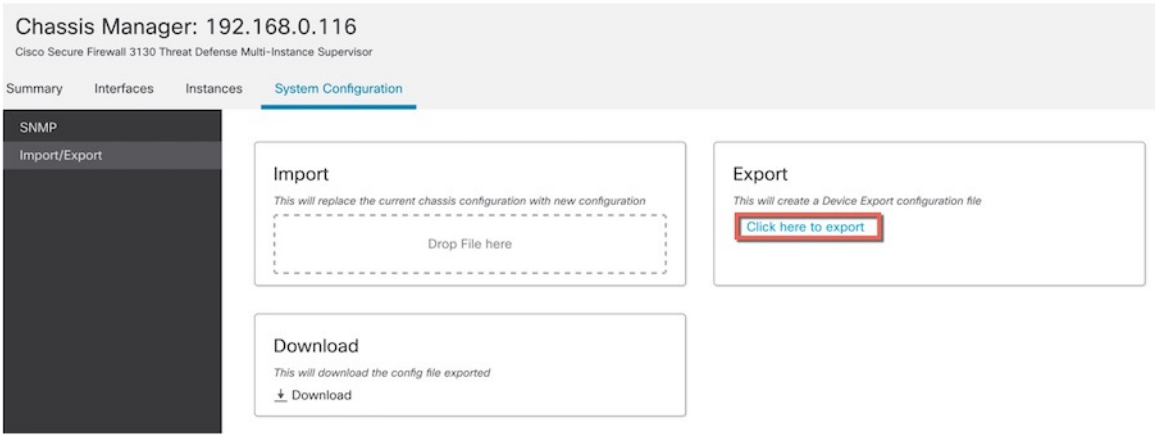
Step 2 Click **System Configuration**.

Step 3 Click **Import/Export**.

Step 4 To export the configuration, follow these steps.

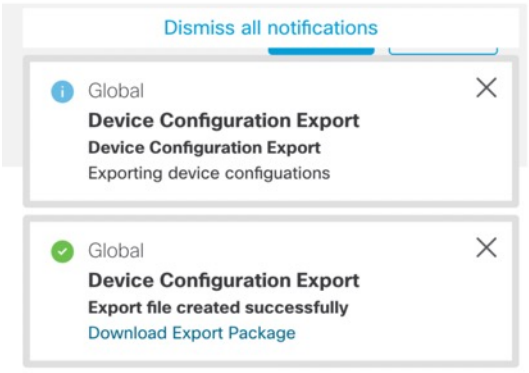
a) In the **Export** area, click **Click here to export**.

Figure 74: Create Export File



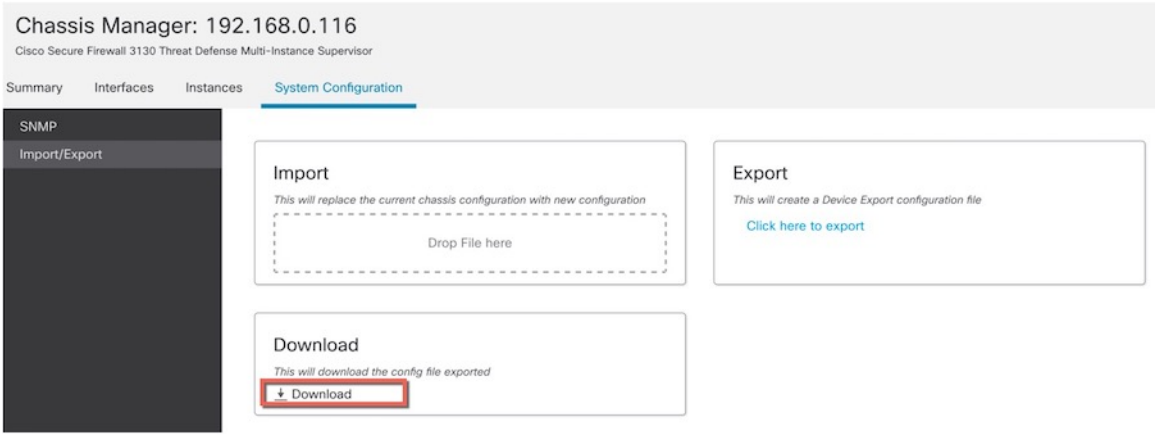
b) Monitor the notifications for the **Export file created successfully** message.

Figure 75: Export File Created Successfully



c) Download the export file by clicking the notification message (**Download Export Package**) or by clicking **Download**.

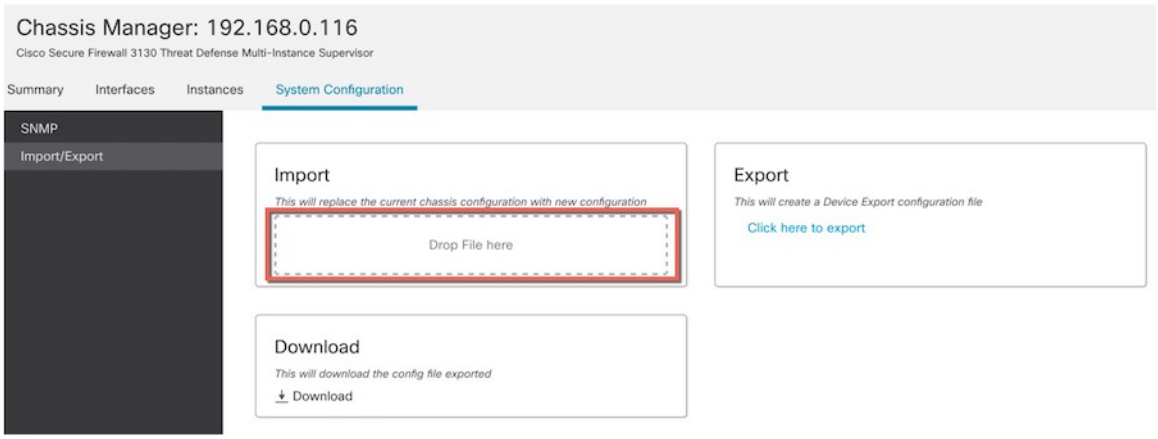
Figure 76: Download



The file is saved with the .sfo extension.

Step 5 To import a configuration, drag the **.sfo** file on the **Import** > **Drop File here** area.

Figure 77: Import



Chassis platform settings

Chassis platform settings are configuration features that configure a range of features for managing the chassis and can be shared among multiple chassis.

Chassis platform settings configure a range of features for managing the chassis. You can share the policy among multiple chassis. If you want different settings per chassis, you must create multiple policies.

Create a chassis platform settings policy

Use the **Platform Settings** page (**Devices** > **Platform Settings**) to manage platform settings policies. This page indicates the type of device for each policy. The **Status** column shows the device targets for the policy.

Procedure

Step 1 Choose **Devices** > **Platform Settings**.

Step 2 For an existing policy, you can **Copy** ( ), **Edit** ( ), or **Delete** ( ) the policy.

Caution
You should not delete a policy that is the last-deployed policy on any of its target devices, even if it is out of date. Before you delete the policy completely, it is good practice to deploy a different policy to those targets.

Step 3 To create a new policy, click **New Policy**.

a) Choose **Chassis Platform Settings** from the drop-down list.
b) Enter a **Name** for the new policy and optionally, a **Description**.
c) Optionally, choose the **Available Chassis** where you want to apply the policy and click **Add** to add the selected chassis.

d) Click **Save**.

The system creates the policy and opens it for editing.

Step 4 To change the target chassis for a policy, click **Edit** ( ) next to the platform settings policy that you want to edit.

a) Click **Policy Assignment**.
b) To assign a chassis to the policy, select it in the **Available Chassis** list and click **Add**.
c) To remove a chassis assignment, click **Delete** ( ) next to a chassis in the **Selected Chassis** list.
d) Click **OK**.

Configure DNS

You need to specify a DNS server if the chassis requires resolution of hostnames to IP addresses. These chassis DNS settings are separate from the DNS settings per instance, which are configured in the device platform settings.

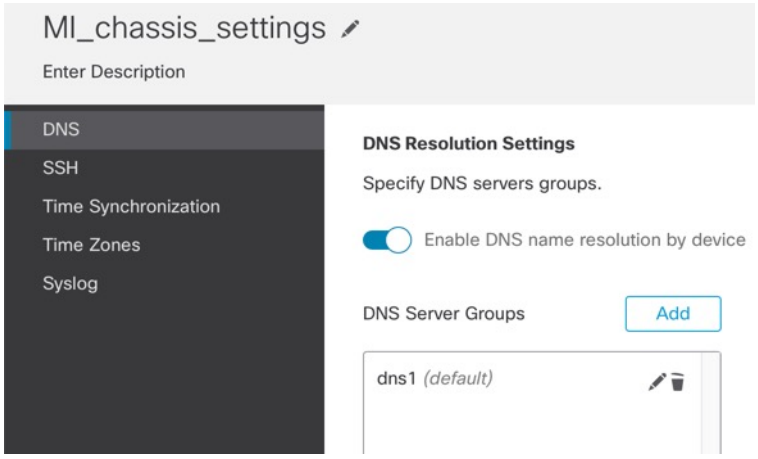
When you configure multiple DNS servers, the chassis uses servers in a random order. You can configure up to four servers across four DNS server groups. For example, you can configure a single server group with four servers, or you can configure four server groups with one server each.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **DNS**.

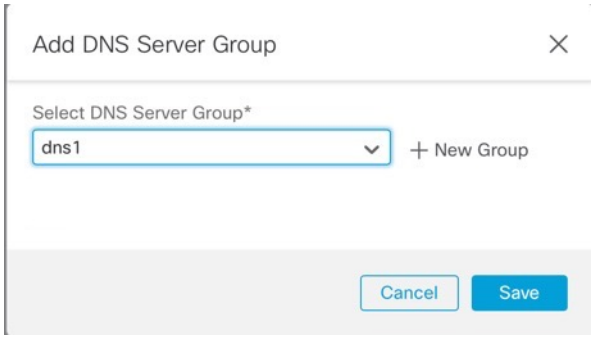
Figure 78: DNS



Step 3 Enable the **Enable DNS name resolution by device** slider.

Step 4 Click **Add** to add a DNS server group.

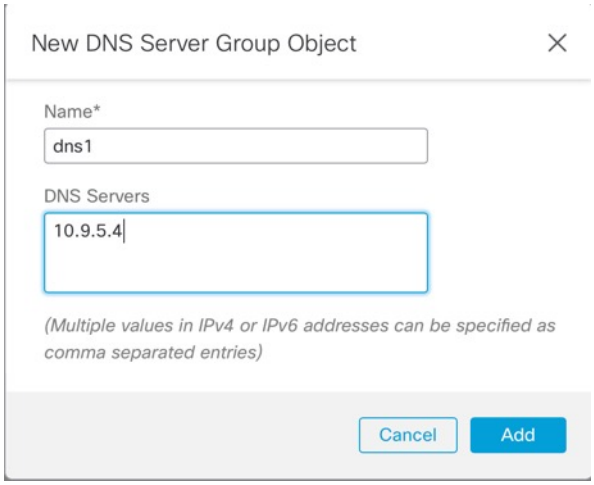
Figure 79: Add DNS server group



Step 5 Either select an existing DNS server group (refer to Create a DNS server group object, on page 2965), or click (+) **New Group**.

If you add a new group, you see this dialog box. Provide a name and up to four DNS server IP addresses as comma-separated values, and click **Add**.

Figure 80: New DNS server group object



Step 6 Click **Save** to add the DNS server to the list.

Step 7 Repeat these steps to add additional server groups.

Make sure you only identify a maximum of four DNS servers in all groups combined.

Step 8 Click **Save** to save all policy changes.

Configure SSH and SSH access list

To allow SSH sessions from the admin user to the chassis on the Management interface, enable the SSH server and configure the allowed networks.

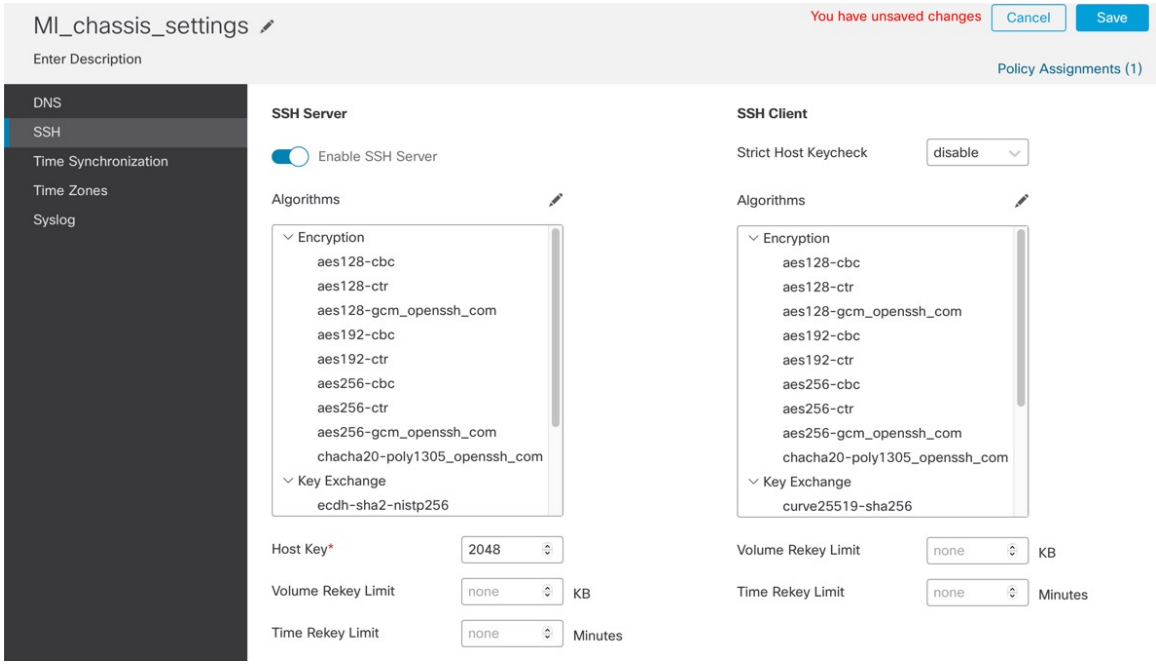
Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **SSH**.

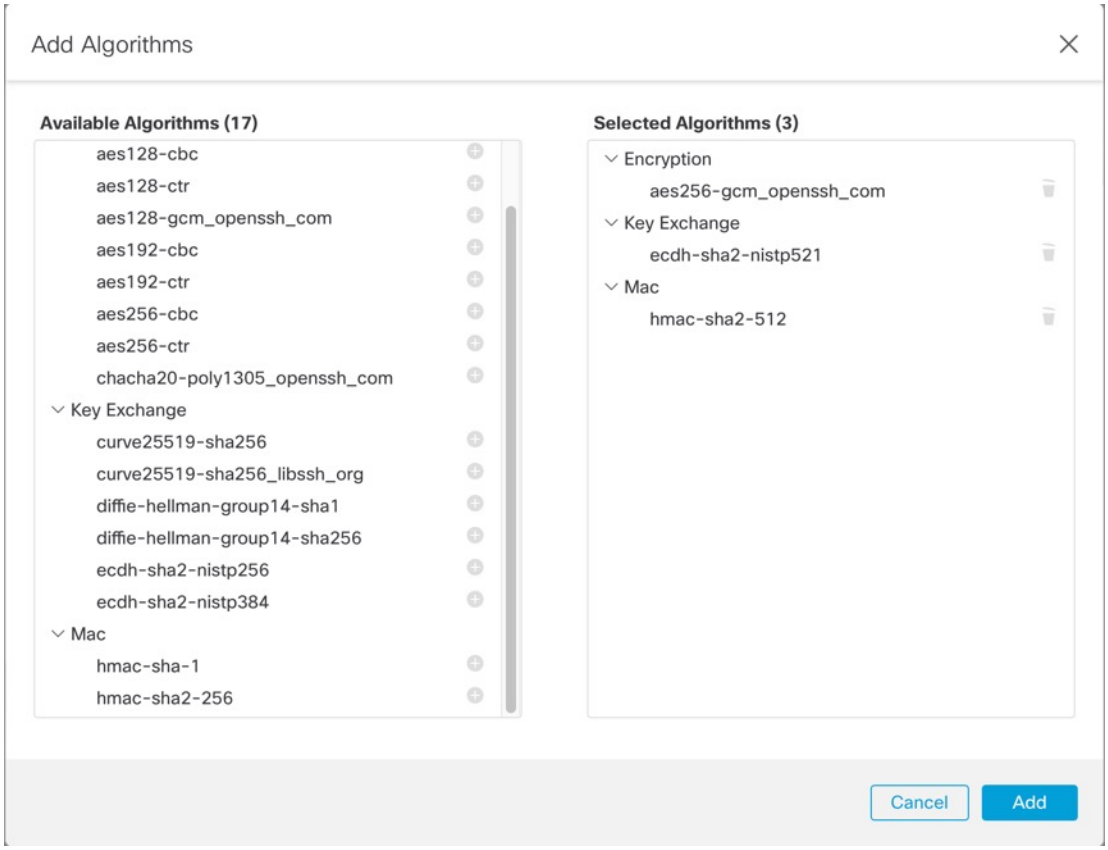
Step 3 To enable SSH access to the chassis, enable the **Enable SSH Server** slider.

Figure 81: SSH



Step 4 To set the allowed **Algorithms**, click **Edit** ( ).

Figure 82: Add Algorithms



a) Select the **Encryption** algorithms.
b) Select the **Key Exchange** algorithms.

   The key exchange provides a shared secret that cannot be determined by either party alone. The key exchange is combined with a signature and the host key to provide host authentication. This key-exchange method provides explicit server authentication.

c) Select the **Mac** integrity algorithms.

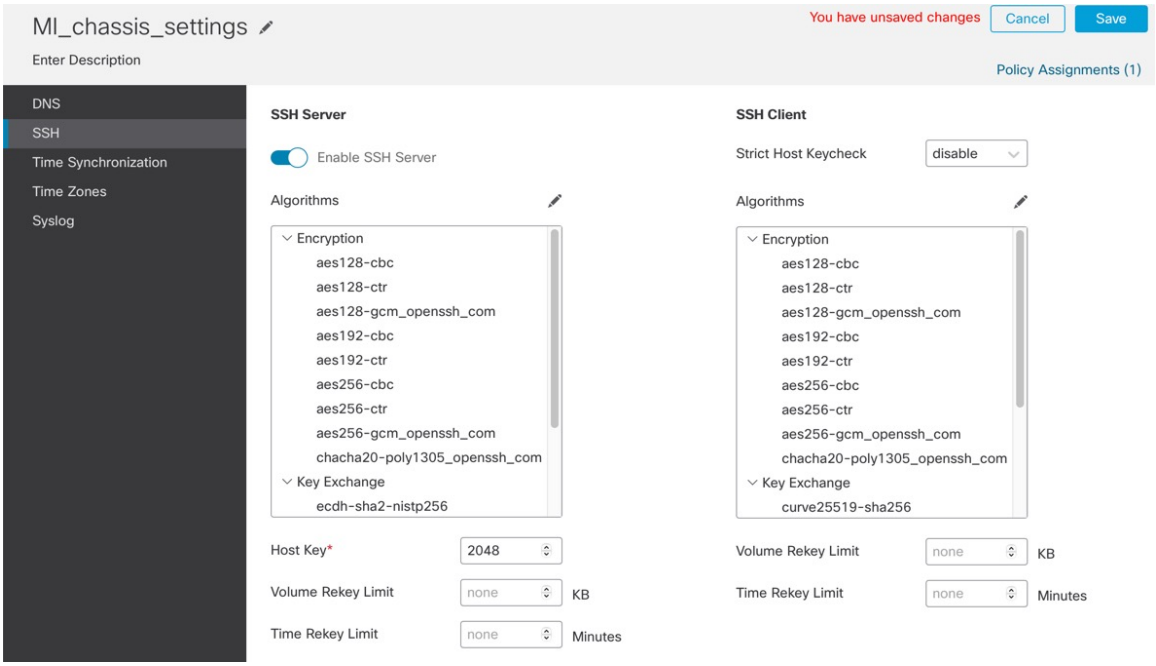
Step 5 For **Host Key**, enter the modulus size for the RSA key pairs.

The modulus value (in bits) is in multiples of 8 from 1024 to 2048. The larger the key modulus size you specify, the longer it takes to generate an RSA key pair. We recommend a value of 2048.

Step 6 For the server **Volume Rekey Limit**, set the amount of traffic in KB allowed over the connection before FXOS disconnects from the session.. For the server **Time Rekey Limit**, set the minutes for how long an SSH session can be idle before FXOS disconnects the session.

Step 7 For the **SSH Client**, configure the following settings.

Figure 83: SSH

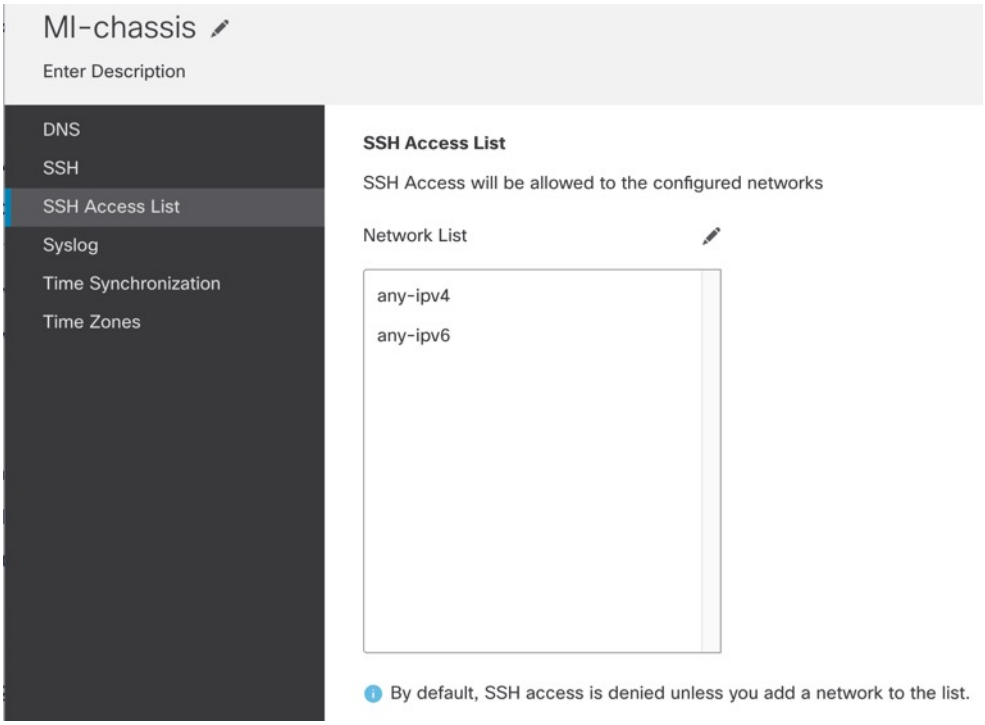


- **Strict Host Keycheck**—Choose **enable**, **disable**, or **prompt** to control SSH host key checking.
  - **enable**—The connection is rejected if the host key is not already in the FXOS known hosts file. You must manually add hosts at the FXOS CLI using the **enter ssh-host** command in the system/services scope.
  - **prompt**—You are prompted to accept or reject the host key if it is not already stored on the chassis.
  - **disable**—(The default) The chassis accepts the host key automatically if it was not stored before.
- **Algorithms**: Click **Edit** (✎). and select the **Encryption**, **Key Exchange**, and **Mac** algorithms.
- **Volume Rekey Limit**: Set the amount of traffic in KB allowed over the connection before FXOS disconnects from the session.
- **Time Rekey Limit**: Set the minutes for how long an SSH session can be idle before FXOS disconnects the session.

Step 8 Choose **SSH Access List**. You need to allow access to IP addresses or networks before you can use SSH.

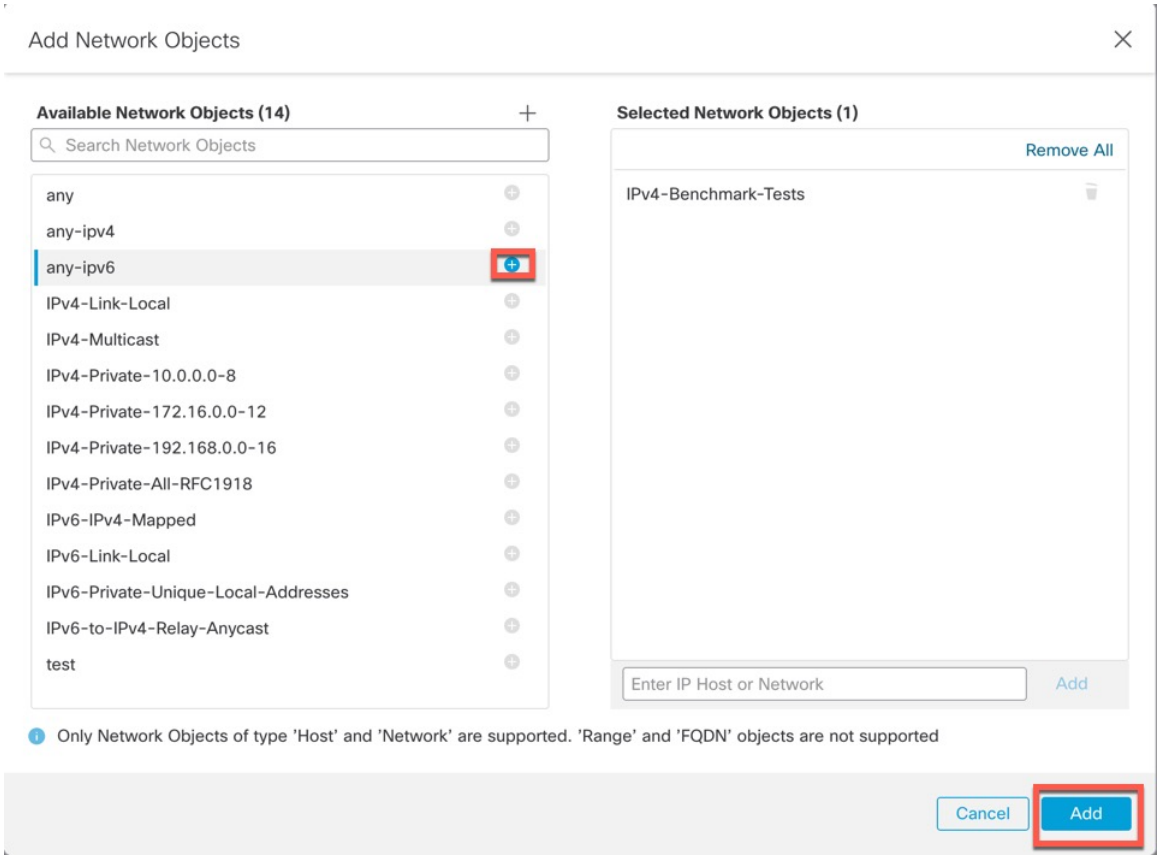
You can add up to 25 access lists.

Figure 84: SSH Access List



Step 9 Click **Edit** ( ) to add network objects and click **Save**. You can also manually enter IP addresses. Click **Save** to save all policy changes.

Figure 85: Network Objects



Configure syslog

You can enable syslogs from the chassis. These syslogs come from the chassis' FXOS operating system.

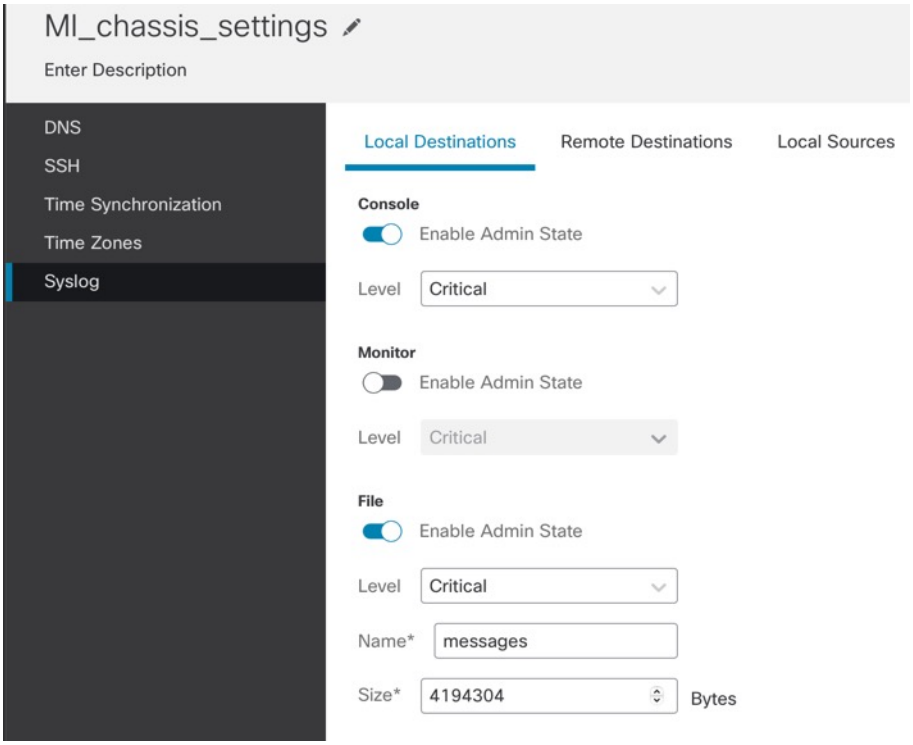
Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **Syslog**.

Step 3 Click the **Local Destinations** tab, and complete the following fields.

Figure 86: Syslog local destinations

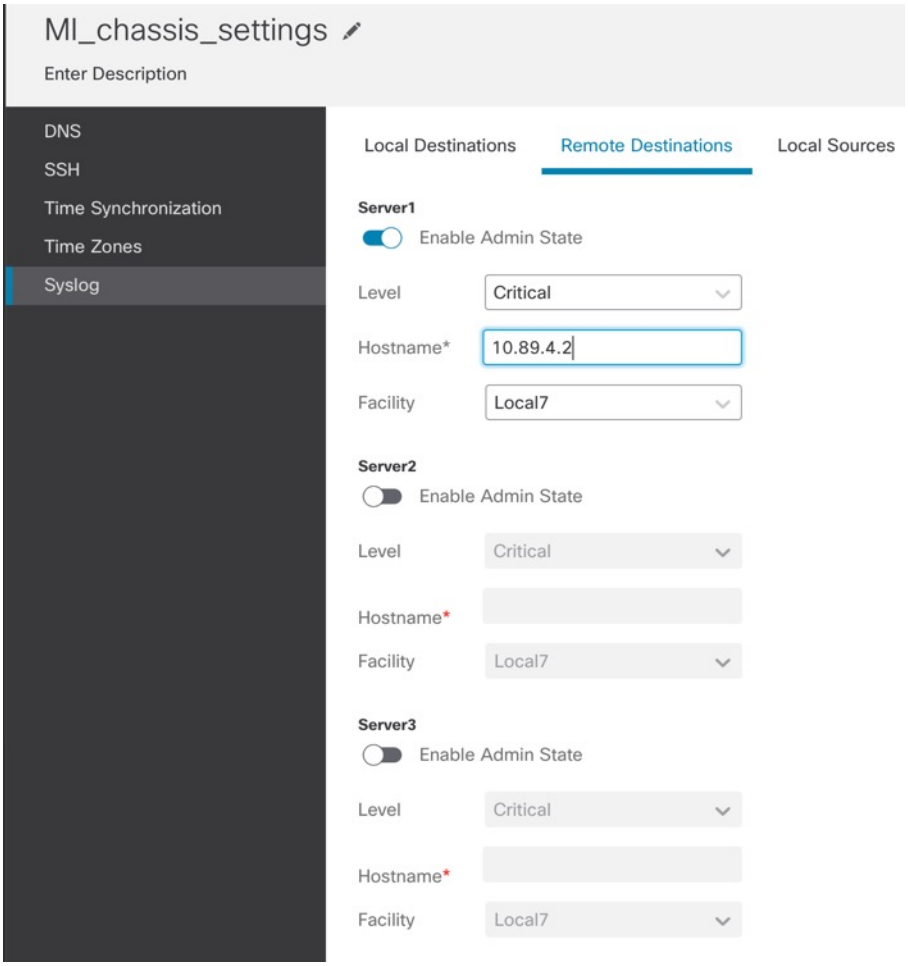


Name	Description
Console Section	
Admin State field	Whether the chassis displays syslog messages on the console. Check the **Enable** check box if you want to have syslog messages displayed on the console as well as added to the log. If the **Enable** check box is unchecked, syslog messages are added to the log but are not displayed on the console.
Level field	If you checked the **Enable** check box for **Console - Admin State**, select the lowest message level that you want displayed on the console. The chassis displays that level and above on the console. This can be one of the following: • **Emergencies** • **Alerts** • **Critical**
Monitor Section	

Name	Description
Admin State field	Whether the chassis displays syslog messages on the monitor. Check the **Enable** check box if you want to have syslog messages displayed on the monitor as well as added to the log. If the **Enable** check box is unchecked, syslog messages are added to the log but are not displayed on the monitor.
Level drop-down list	If you checked the **Enable** check box for **Monitor - Admin State**, select the lowest message level that you want displayed on the monitor. The system displays that level and above on the monitor. This can be one of the following: • **Emergencies** • **Alerts** • **Critical** • **Errors** • **Warnings** • **Notifications** • **Information** • **Debugging**

Step 4 On the **Remote Destinations** tab, complete the following fields for up to three external logs that can store messages generated by the chassis:

Figure 87: Syslog remote destinations



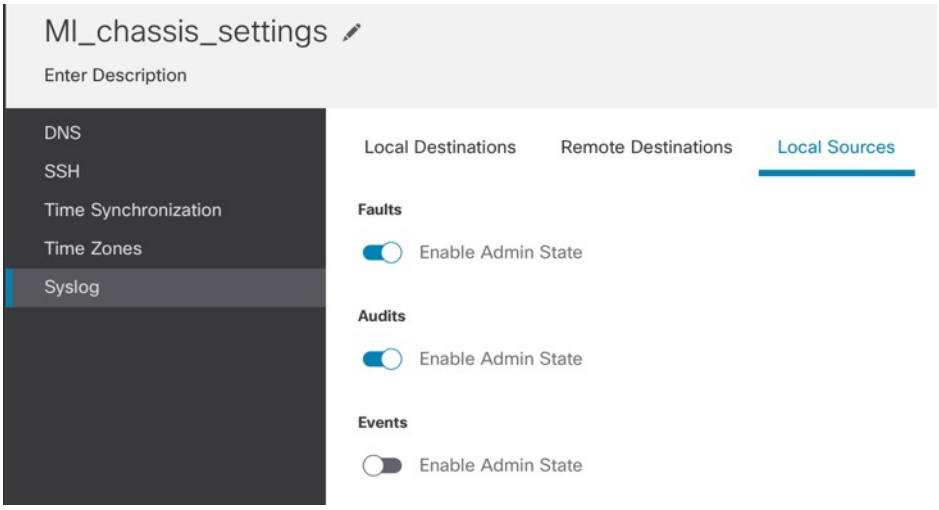
By sending syslog messages to a remote destination, you can archive messages according to the available disk space on the external syslog server, and manipulate logging data after it is saved. For example, you could specify actions to be executed when certain types of syslog messages are logged, extract data from the log and save the records to another file for reporting, or track statistics using a site-specific script.

Name	Description
Admin State field	Check the **Enable** check box if you want to have syslog messages stored in a remote log file.

Name	Description
Level drop-down list	Select the lowest message level that you want the system to store. The system stores that level and above in the remote file. This can be one of the following: • **Emergencies** • **Alerts** • **Critical** • **Errors** • **Warnings** • **Notifications** • **Information** • **Debugging**
Hostname/IP Address field	The hostname or IP address on which the remote log file resides. **Note** You must configure a DNS server if you use a hostname rather than an IP address.
Facility drop-down list	Choose a system log facility for syslog servers to use as a basis to file messages. This can be one of the following: • **Local0** • **Local1** • **Local2** • **Local3** • **Local4** • **Local5** • **Local6** • **Local7**

Step 5 Click the **Local Sources** tab, and complete the following fields.

Figure 88: Syslog local sources



Name	Description
Faults > **Enable Admin State**	Enable system fault logging.
Audits > **Enable Admin State**	Enable audit logging.
Events > **Enable Admin State**	Enable system event logging.

Step 6 Click **Save** to save all policy changes.

Configure time synchronization

NTP is used to implement a hierarchical system of servers that provide precisely synchronized time among network systems. This kind of accuracy is required for time-sensitive operations, such as validating CRLs, which include a precise time stamp. You can configure up to four NTP servers.

Note

- FXOS uses NTP version 3.
- If the stratum value of an external NTP server is 13 or greater, the application instance cannot sync to the NTP server on the FXOS chassis. Each time a NTP client syncs to a NTP server, the stratum value increases by one.

  If you have set up your own NTP server, you can find its stratum value in the /etc/ntp.conf file on the server. If the NTP server has stratum value of 13 or greater you can either change the stratum value in the ntp.conf file and restart the server, or use a different NTP server (for example: pool.ntp.org).

Before you begin

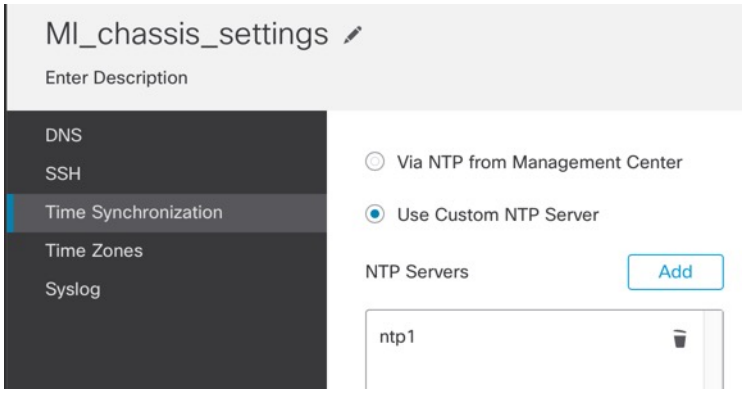
If you use a hostname for the NTP server, you must configure a DNS server. For more information, refer to Configure DNS, on page 653.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **Time Synchronization**.

Figure 89: Time synchronization



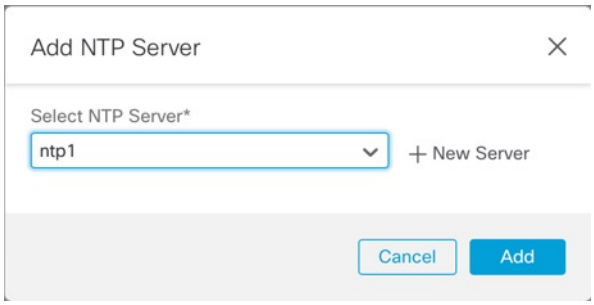
Step 3 If you want to obtain the time from the Cloud-Delivered Firewall Management Center, click **Via NTP from Management Center**.

This option ensures both the chassis and the Cloud-Delivered Firewall Management Center have the same time.

Step 4 To use an external NTP server, click **Use Custom NTP Server**.

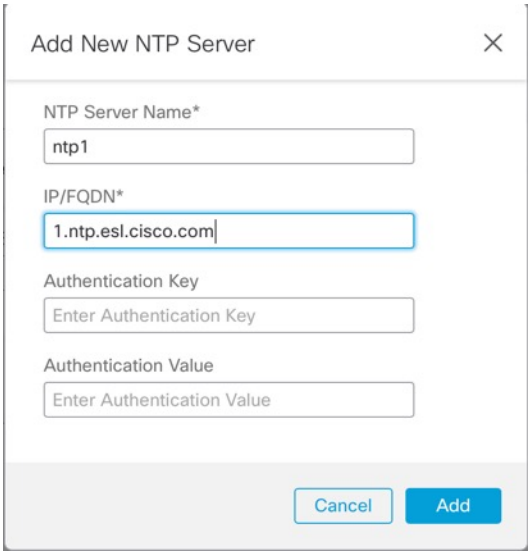
a) Click **Add** to add a server.

Figure 90: Add NTP server



b) Choose any already-defined servers from the drop-down menu and click **Add**, or click + **New Server** to add a new server.

Figure 91: Add new NTP server



c) For a new server, enter the following fields, and click **Add**.

- **NTP Server Name**: A name to identify this server.
- **IP/FQDN**: The IP address or hostname of the server.
- **Authentication Key** and **Authentication Value**: Obtain the key ID and value from the NTP server. For example, to generate the SHA1 key on NTP server Version 4.2.8p8 or later with OpenSSL installed, enter the **NTP-keygen -M** command, and then view the key ID and value in the NTP.keys file. The key is used to tell both the client and server which value to use when computing the message digest.

  Only SHA1 is supported for NTP server authentication.

Step 5 Click **Save** to save all policy changes.

Configure time zones

Time zone configuration is essential for maintaining accurate timestamps in system logs and ensuring proper time-based functionality across the device.

Before you begin

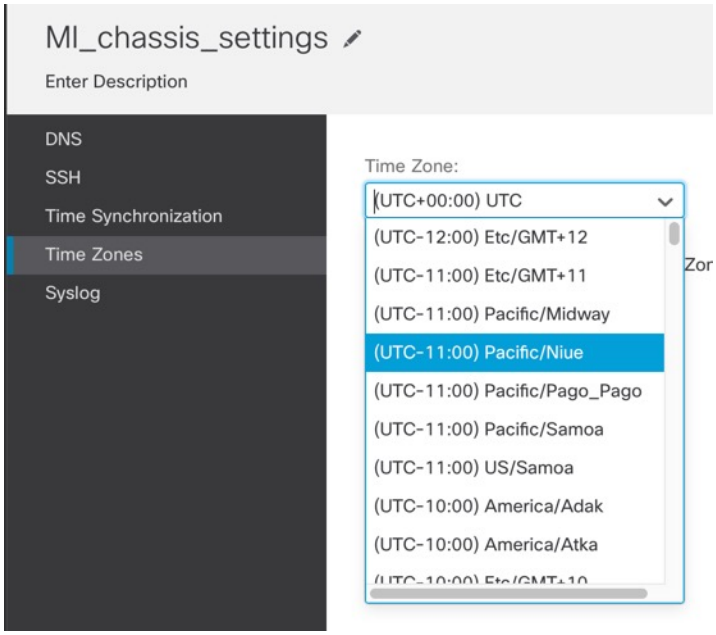
Follow these steps to configure the time zone for the chassis:

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **Time Zones**.

Figure 92: Time zones



Step 3 Choose your **Time Zone** from the drop-down menu.

Step 4 Click **Save** to save all policy changes.

Multi-instance mode management

This section describes less common tasks, including changing settings at the FXOS CLI or changing interfaces assigned to the chassis.

Onboard the multi-instance chassis using the CLI

If you want to pre-configure your devices for multi-instance mode before you add them to the Cloud-Delivered Firewall Management Center, you can perform this procedure. To use the Cloud-Delivered Firewall Management Center to convert to multi-instance mode, refer to Convert a device to Multi-Instance mode, on page 629.

You need to connect to the chassis FXOS CLI at the console port to enable multi-instance mode and perform initial setup. After you configure the chassis, you can add it to the Cloud-Delivered Firewall Management Center. The management center and the chassis share a separate management connection using the chassis MGMT interface.

Note Although you can connect to SSH on the management port, we recommend using the console port to avoid multiple disconnections. If you use SSH before you switch to multi-instance mode, you connect to the Firewall Threat Defense CLI and need to connect to FXOS using **connect FXOS**. This procedure covers the console port.

Use the Cloud-Delivered Firewall Management Center to configure all chassis settings as well as instances. The Secure Firewall Chassis Manager or configuration at the FXOS CLI is not supported.

Procedure

Step 1 Connect to the chassis console port.

The console port connects to the FXOS CLI.

Step 2 Log in with the username **admin** and the password **Admin123**.

The first time you log in to FXOS, you are prompted to change the password.

Note
If the password was already changed, and you do not know it, you must reimage the device to reset the password to the default. See the FXOS troubleshooting guide for the reimage procedure.

Example:

```
firepower login: admin
Password: Admin123
Successful login attempts for user 'admin' : 1

[...]

Hello admin. You must change your password.
Enter new password: ********
Confirm new password: ********
Your password was updated successfully.

[...]

firepower#
```

Step 3 Check your current mode, Native or Container. If the mode is Native, you can continue with this procedure to convert to multi-instance (Container) mode.

show system detail

Example:

```
firepower # show system detail

Systems:
    Name: firepower
    Mode: Stand Alone
    System IP Address: 172.16.0.50
    System IPv6 Address: ::
    System Owner:
    System Site:
    Deploy Mode: Native
    Description for System:
firepower #
```

Step 4 Enable multi-instance mode.

scope system

set deploymode container

You are prompted to reboot.

Example:

```
firepower# scope system
firepower /system # set deploymode container
All configuration and bootable images will be lost and system will reboot.
If there was out of band upgrade, it might reboot with the base version and need to re-image
 to get the expected running version.
Do you still want to change deploy mode? (yes/no):yes
firepower /system #
```

To change the mode back to appliance mode, enter **set deploymode native**.

Step 5 After the chassis finishes rebooting, set the Management IP address. You can use IPv4 and/or IPv6. The default address is 192.168.45.45/24 with a gateway of 192.168.45.1.

IPv4:

scope fabric-interconnect

set out-of-band static ip *ip_address* **netmask** *network_mask* **gw** *gateway_ip_address*

IPv6:

scope fabric-interconnect

scope ipv6-config

set out-of-band static ipv6 *ipv6_address* **ipv6-prefix** *prefix_length* **ipv6-gw** *gateway_address*

Example:

IPv4:

```
firepower-3110# scope fabric-interconnect
firepower-3110 /fabric-interconnect # set out-of-band static ip 10.5.23.8 netmask
255.255.255.0 gw 10.5.23.1
```

IPv6:

```
firepower-3110# scope fabric-interconnect
firepower-3110 / fabric-interconnect # scope ipv6-config
firepower-3110 / fabric-interconnect /ipv6-config # set out-of-band static ipv6 2001:DB8::34
 ipv6-prefix 64 ipv6-gw 2001:DB8::1
```

Step 6 Log in to Security Cloud Control.

Step 7 In the left pane, click **Security Devices**.

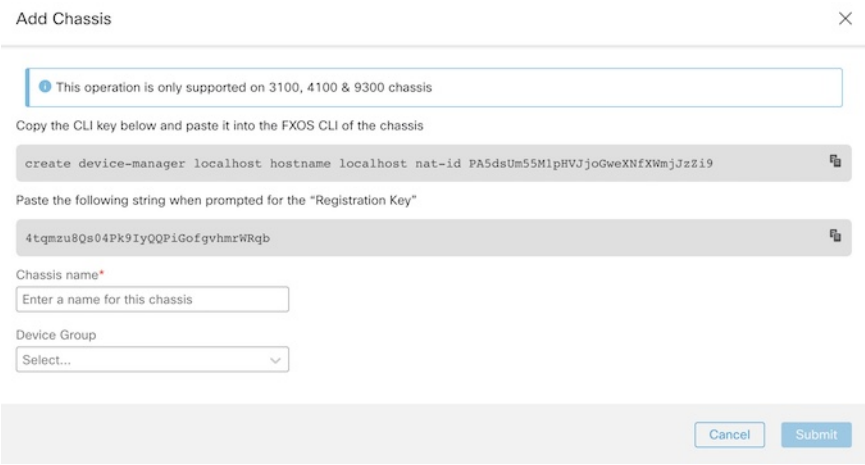
Step 8 Click **Onboard device or service** ( ) icon.

Step 9 Click the **FTD Chassis** tile to open the **Add Chassis** dialog box.

Figure 93: FTD Chassis Tile



Figure 94: Add Chassis



Step 10 Click **Copy** ( ) to copy the top generated command, then paste it at the FXOS CLI of your chassis.

Step 11 When prompted for the Registration Key at the FXOS CLI, click **Copy** ( ) on the **Add Chassis** dialog box for the generated registration key and paste it at the FXOS CLI.

You can disconnect from the FXOS CLI at this point.

Step 12 In the Cloud-Delivered Firewall Management Center **Chassis Name** field, enter a name for the chassis as you want it to display in the Cloud-Delivered Firewall Management Center.

Step 13 (Optional) Add the chassis to a **Device Group**.

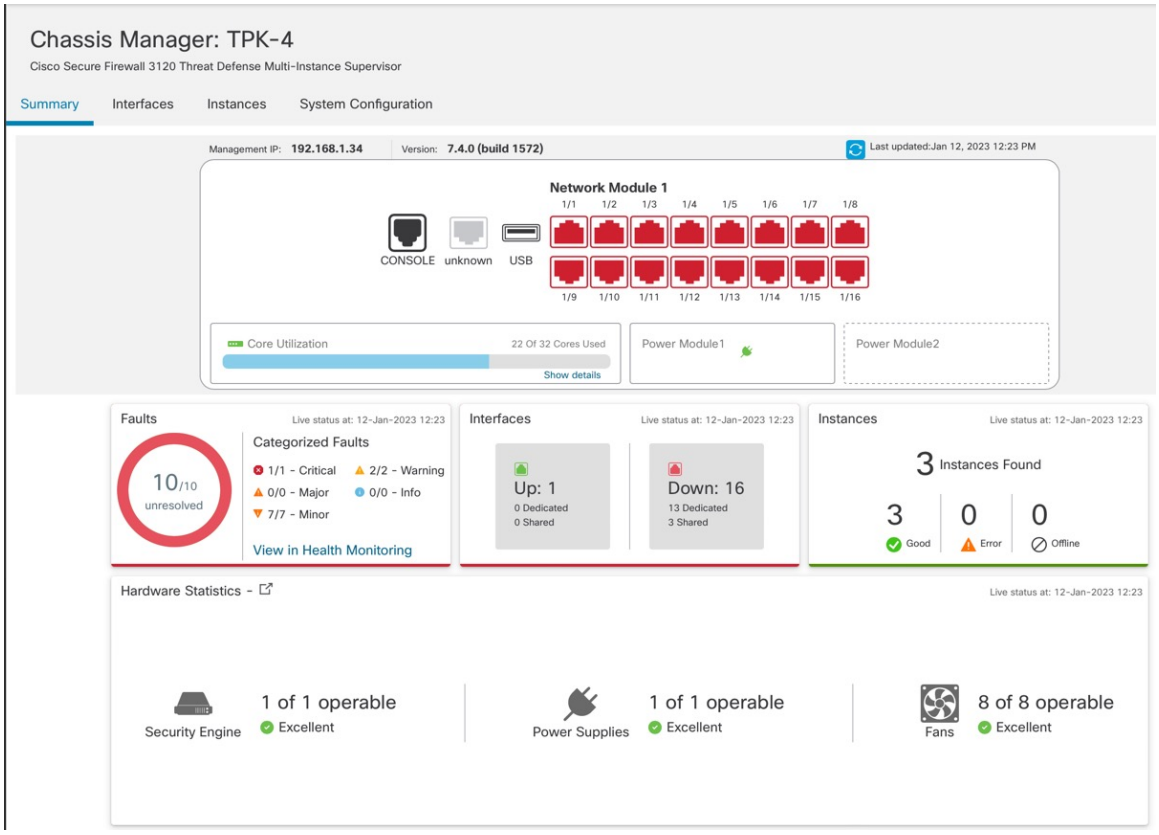
Step 14 Click **Submit**.

The chassis is added to the **Device** > **Device Management** page.

Step 15 To view and configure the chassis, click **Manage** in the **Chassis** column, or click **Edit** ( ) .

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Figure 95: Chassis Summary



Change interfaces assigned to an instance

You can allocate or unallocate an interface on the instance. Adding a new interface, or deleting an unused interface, has minimal impact on the instance configuration. You can also edit the membership of an allocated EtherChannel without affecting the instance.

Deleting an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the instance configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface.

Policies that refer to security zones are not affected.

Note For high availability, you need to make the same interface changes for the other unit. Otherwise, high availability might not operate correctly.

Before you begin

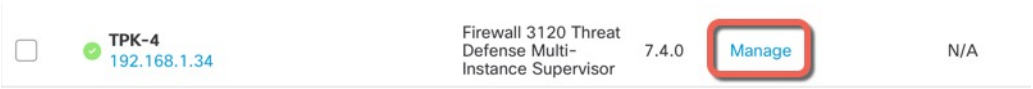
- Configure your interfaces according to Instance configuration, on page 629.

- If you plan to add an already-allocated interface to an EtherChannel, you need to unallocate the interface from the instance first, then add the interface to the EtherChannel. For a new EtherChannel, you can then allocate the EtherChannel to the instance.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

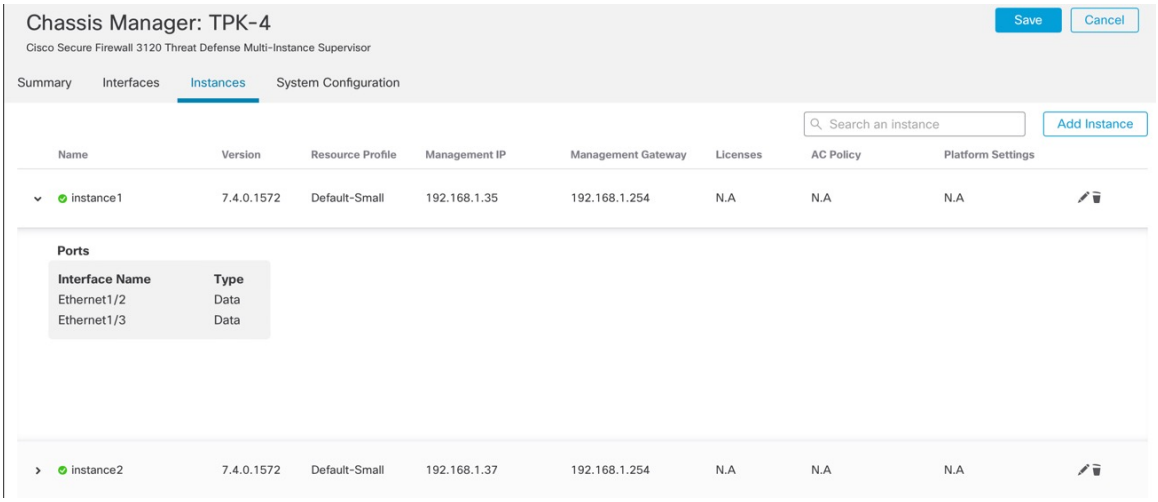
Figure 96: Manage Chassis



The **Chassis Manager** page opens for the chassis to the **Summary** page.

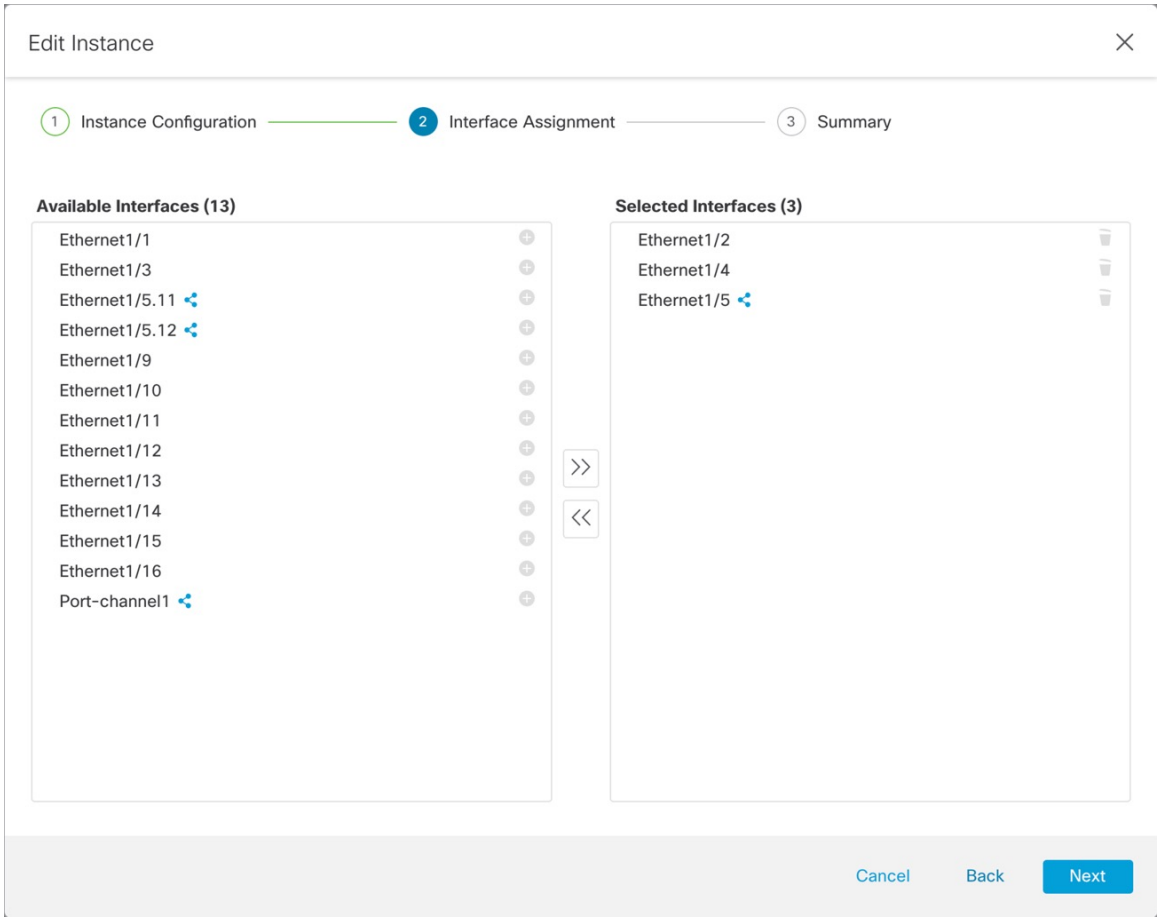
Step 2 Click **Instances**, and click **Edit** ( ) next to the instance for which you want to change interfaces.

Figure 97: Instances



Step 3 Click **Next** until you get to the **Interface Assignment** screen.

Figure 98: Interface assignment



Shared interfaces show the sharing icon ( ).

Step 4 Make your interface changes, and then click **Next**.

Step 5 Click **Save** on the **Summary** screen.

Step 6 For high availability, you need to make the same interface changes for the other unit. Otherwise, high availability might not operate correctly.

Change chassis management settings at the FXOS CLI

If you want to change the chassis management interface IP address and gateway, change the Cloud-Delivered Firewall Management Center to a new manager, change the admin password, or disable multi-instance mode, you can do so from the FXOS CLI.

Procedure

Step 1 Connect to the chassis console port.

The console port connects to the FXOS CLI.

Note
We recommend using the console port. You can also connect using SSH to the management interface, if configured in the chassis platform settings in the Cloud-Delivered Firewall Management Center; however, if you change the management IP address, you will be disconnected.

Step 2 Log in with the username **admin** and the password you set during initial setup.

Step 3 Change the Management IP address. You can use a static IPv4 and/or IPv6 address.

IPv4:

scope fabric-interconnect

set out-of-band static IP *ip_address* **netmask** *network_mask* **gw** *gateway_ip_address*

IPv6:

scope fabric-interconnect

scope ipv6-config

set out-of-band static ipv6 *ipv6_address* **ipv6-prefix** *prefix_length* **ipv6-gw** *gateway_address*

Example:

IPv4:

```
firepower-3110# scope fabric-interconnect
firepower-3110 /fabric-interconnect # set out-of-band static ip 10.5.23.8 netmask
255.255.255.0
gw 10.5.23.1
```

IPv6:

```
firepower-3110# scope fabric-interconnect
firepower-3110 / fabric-interconnect # scope ipv6-config
firepower-3110 / fabric-interconnect /ipv6-config # set out-of-band static ipv6 2001:DB8::34
ipv6-prefix 64 ipv6-gw 2001:DB8::1
```

Step 4 Change the Cloud-Delivered Firewall Management Center.

You should first unregister the chassis from the current Cloud-Delivered Firewall Management Center.

enter device-manager *manager_name* [**hostname** {*hostname* | *ipv4_address* | *ipv6_address*}] [**NAT-ID** *nat_id*]

You are prompted for the registration key.

You can enter this command from any scope.

- **hostname** {*hostname* | *ipv4_address* | *ipv6_address*}: Specifies either the FQDN or IP address of the Cloud-Delivered Firewall Management Center. At least one of the devices, either the Cloud-Delivered Firewall Management Center or the chassis, must have a reachable IP address to establish the two-way, TLS-1.3-encrypted communication channel between the two devices. If you do not specify a **hostname**, then the chassis must have a reachable IP address or hostname and you must specify the **NAT-ID**.

- **NAT-ID** *nat_id*: Specifies a unique, one-time string of your choice that you will also specify on the Cloud-Delivered Firewall Management Center when you register the chassis when one side does not specify a reachable IP address or hostname. It is required if you do not specify a **hostname**, however we recommend that you always set the NAT ID even when you specify a hostname or IP address. The NAT ID must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID cannot be used for any other devices registering to the Cloud-Delivered Firewall Management Center.
- **Registration Key:** *reg_key*: You will be prompted for a one-time registration key of your choice that you will also specify on the Cloud-Delivered Firewall Management Center when you register the chassis. The registration key must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-).

Example:

```
firepower-3110# enter device-manager boulder_fmc hostname 10.89.5.35 nat-id 93002
(Valid registration key characters: [a-z],[A-Z],[0-9],[-]. Length: [2-36])
Registration Key: Impala67
```

Step 5 Change the admin password.

scope security

set password

Enter a password: *password*

Confirm the password: *password*

Example:

```
firepower-3110# scope security
firepower-3110 /security # set password
Enter new password: Sw@nsong67
Confirm new password: Sw@nsong67
firepower-3110 /security #
```

Step 6 Disable multi-instance mode and set the system back to appliance mode.

scope system

set deploymode native

You are prompted to reboot.

Example:

```
firepower-3110# scope system
firepower-3110 /system # set deploymode native
All configuration and bootable images will be lost and system will reboot.
If there was out of band upgrade, it might reboot with the base version and
need to re-image to get the expected running version.
Do you still want to change deploy mode? (yes/no):yes
firepower-3110 /system #
```

To change the mode back to multi-instance mode, enter **set deploymode container**. You can check the current mode using the **show system detail** command.

Multi-instance mode

This section helps you troubleshoot and diagnose your multi-instance mode chassis and instances.

Multi-instance setups

A multi-instance setup is a system configuration approach that

- enables running multiple independent instances on a single hardware platform,
- allows flexible allocation of resources to each instance, and
- supports both container and native (appliance) operating modes.

show system detail

This FXOS command shows the current mode, Native or Container. If the mode is Native (also known as appliance mode), you can convert to multi-instance (Container) mode. Note that the prompt/name in multi-instance mode is generic "firepower-<model>" while the prompt in appliance mode is the hostname you set for the Firewall Threat Defense (by default "firepower")

```
firepower # show system detail

Systems:
    Name: firepower
    Mode: Stand Alone
    System IP Address: 172.16.0.50
    System IPv6 Address: ::
    System Owner:
    System Site:
    Deploy Mode: Native
    Description for System:
firepower #
```

scope system > show

This FXOS command shows the current mode in a table format. Note that the prompt/name in multi-instance mode is generic "firepower-<model>" while the prompt in appliance mode is the hostname you set for the Firewall Threat Defense.

```
firepower-3110# scope system
firepower-3110 /system # show

Systems:
    Name       Mode        Deploy Mode System IP Address System IPv6 Address
    ---------- ----------- ----------- ----------------- -------------------
```

```
    firepower-3110
               Stand Alone Container   10.89.5.42        ::

3110-1# scope system
3110-1 /system # show

Systems:
    Name       Mode        Deploy Mode System IP Address System IPv6 Address
    ---------- ----------- ----------- ----------------- -------------------
    3110-1     Stand Alone Native      10.89.5.41        ::
3110-1 /system #
```

Instance interfaces

An instance interface is a physical or virtual network interface on a device that

- enables connectivity between multiple logical instances within a system,
- facilitates assignment of interfaces to one or more instances, and
- supports configuration of forwarding rules using ECMP and multicast groups.

show portmanager switch forward-rules hardware mac-filter

This command shows the internal switch-forwarding rule for two instances with a dedicated physical interface assigned to each instance. Ethernet 1/2 is assigned to ftd1 and Ethernet 1/1 is assigned to ftd2.

ECMP group 1540 is assigned to ftd1 and ECMP group 1541 is assigned to ftd2.

```
secfw-3140(local-mgmt)# show portmanager switch forward-rules hardware mac-filter
      VLAN   SRC_PORT   PC_ID   SRC_ID   DST_PORT     PKT_CNT      DMAC
1        0         17       0       17         19       29164  0:0:0:0:0:0
2        0         19       0       19         17       67588  0:0:0:0:0:0
3        0          1       0      101       1541           0  a2:5b:83:0:0:15
4        0          1       0      101       1541        8181  ff:ff:ff:ff:ff:ff
5        0          2       0      102       1540           0  a2:5b:83:0:0:18
6        0          2       0      102       1540         431  ff:ff:ff:ff:ff:ff
7        0         17       0        0          0       11133  0:0:0:0:0:0
8        0         17       0        0          0           0  0:0:0:0:0:0
```

This command shows the internal switch-forwarding rule for two instances with shared physical interfaces assigned to two instances. Ethernet 1/1 is shared between ftd1 and ftd2.

ECMP group 1540 is assigned to ftd1 and ECMP group 1541 is assigned to ftd2.

MCAST group 4096 is used for replicating broadcast traffic between ftd1 and ftd2.

```
firepower-3140(local-mgmt)# show portmanager switch forward-rules hardware mac-filter
      VLAN   SRC_PORT   PC_ID   SRC_ID   DST_PORT     PKT_CNT      DMAC
1        0         17       0       17         19        2268  0:0:0:0:0:0
2        0         19       0       19         17        4844  0:0:0:0:0:0
3        0          1       0      101       1541           0  a2:5b:83:0:0:9
4        0          1       0      101       4096         546  ff:ff:ff:ff:ff:ff
5        0          1       0      101       1540           0  a2:5b:83:0:0:c
6        0         17       0        0          0        1263  0:0:0:0:0:0
7        0         17       0        0          0           0  0:0:0:0:0:0
```

This command shows the internal switch-forwarding rule for two instances with shared subinterfaces assigned to both instances. Ethernet 1/1.2452 is shared between ftd1 and ftd2.

ECMP group 1540 is assigned to ftd1 and ECMP group 1541 is assigned to ftd2.

MCAST group 4097 is used for replicating broadcast traffic between ftd1 and ftd2.

```
firepower-3140(local-mgmt)# show portmanager switch forward-rules hardware mac-filter
      VLAN   SRC_PORT   PC_ID   SRC_ID   DST_PORT     PKT_CNT      DMAC
1        0         17       0       17         19       21305  0:0:0:0:0:0
2        0         19       0       19         17       50976  0:0:0:0:0:0
3     2452          1       0      101       1541         430  a2:5b:83:0:0:f
4     2452          1       0      101       4097           0  ff:ff:ff:ff:ff:ff
5     2452          1       0      101       1540           0  a2:5b:83:0:0:12
6        0         17       0        0          0       11038  0:0:0:0:0:0
7        0         17       0        0          0           0  0:0:0:0:0:0
```

show portmanager switch ecmp-groups detail

Use this command to list each Instance Ecmp-Vport-Physical port mapping detail.

Note Physical-Port 18 is the backplane uplink interface between the internal switch and the instance.

```
firepower-3140(local-mgmt)# show portmanager switch ecmp-groups detail
       ECMP-GROUP  VPORT      PHYSICAL-PORT
1      1536        256        18
2      1537        257        18
3      1538        258        18
4      1539        259        18
5      1540        260        18
6      1541        261        18
7      1542        262        18
8      1543        263        18
9      1544        264        18
10     1545        265        18
```

show portmanager switch mcast-groups detail

Use this command to list MCAST group membership details.

```
firepower-3140(local-mgmt)# show portmanager switch mcast-groups detail
       MCAST-GROUP
1      4096
       Member-ports
       Ethernet 1/1
       ECMP-ID 1541
       ECMP-ID 1540
```

show portmanager counters mcast-group

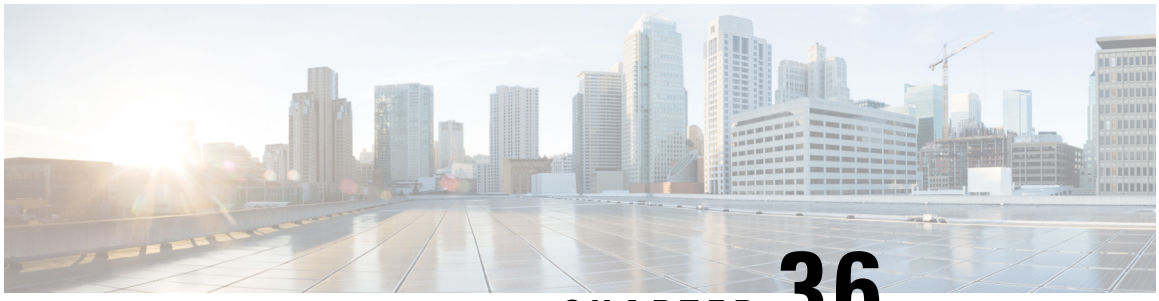
Use this command to check the MCAST group packet counter.

```
firepower-3140(local-mgmt)# show portmanager counters mcast-group 4096
PKT_CNT: 8106
```

show portmanager counters ecmp

Use this command to check the ECMP group packet counter.

```
firepower-3140(local-mgmt)# show portmanager counters ecmp 1541
PKT_CNT: 430
```



CHAPTER 36

High Availability for Devices

The following topics describe how to configure Active/Standby failover to accomplish high availability of the Firewall Threat Defense.

- Secure Firewall Threat Defense high availability, on page 681
- Config-sync optimization, on page 700
- Requirements and prerequisites for high availability, on page 702
- Guidelines for High availability, on page 702
- Add a high availability pair, on page 705
- Configure optional high availability parameters, on page 708
- Manage High availability, on page 710
- Monitoring High availability, on page 717
- High availability break troubleshooting in remote branch deployments, on page 718

Secure Firewall Threat Defense high availability

Secure Firewall Threat Defense high availability is a failover configuration that

- requires two identical Firewall Threat Defense devices connected through a dedicated failover link and optionally a state link
- supports Active/Standby failover where one unit passes traffic while the standby unit does not actively pass traffic synchronizes configuration and state information, and
- monitors the health of the active unit (hardware, interfaces, software, and environmental status) to determine if failover conditions are met.

Failover process details

When a failover occurs, the active unit fails over to the standby unit, which then becomes active. The health of the active unit (hardware, interfaces, software, and environmental status) is monitored to determine if specific failover conditions are met. If those conditions are met, failover occurs.

Note High availability is not supported on Firewall Threat Defense Virtual running in the public cloud. See the Secure Firewall Threat Defense Virtual getting started guides for more information about configuring the Firewall Threat Defense Virtual device for high availability.

High availability support on Firewall Threat Defense devices in a remote branch office deployment

High availability support on Firewall Threat Defense devices in a remote branch office deployment is a feature that

- enables Security Cloud Control to manage devices through data interfaces instead of management interfaces
- supports centralized management through outside Security Cloud Control access for single internet connection environments, and
- requires software version 7.2 or later on managed devices.

Remote branch office deployment characteristics

In a remote branch office deployment, the data interface of the Firewall Threat Defense device is used for Security Cloud Control management instead of the management interface on the device. Most remote branch offices only have a single internet connection. Outside Security Cloud Control access makes centralized management possible.

You can use *any* data interface for Security Cloud Control access, for example, the inside interface if you have an inside Security Cloud Control. However, this guide primarily covers outside interface access, because it is the most likely scenario for remote branch offices.

Security Cloud Control provides high availability support on the Firewall Threat Defense devices that it manages through the data interface. This feature is supported on devices running on software version 7.2 or later.

For more information, see *Firepower Threat Defense Deployment with a Remote FMC* in the Cisco Firepower Getting Started Guide.

High availability system requirements

High availability system requirements are specifications that define the hardware, software, and license prerequisites for Firewall Threat Defense devices to function in a High availability configuration.

Hardware requirements

Hardware requirements are specifications that ensure two units in a High availability configuration can

- maintain identical hardware configurations including model, interfaces, modules, and RAM
- support proper synchronization and failover functionality, and
- ensure reliable high availability operation.

Hardware requirement specifications

The two units in a High availability configuration must meet these requirements:

- Be the same model. In addition, for container instances, they must use the same resource profile attributes.

For the Firepower 9300, High Availability is only supported between same-type modules. However, the two chassis can include mixed modules. For example, each chassis has an SM-56, SM-48, and SM-40. You can create High Availability pairs between the SM-56 modules, between the SM-48 modules, and between the SM-40 modules.

If you change the resource profile after you add the High Availability pair to the Security Cloud Control, update the inventory for each unit on the **Devices** > **Device Management** > **Device** > **System** > **Inventory** dialog box.

If you assign a different profile to instances in an established high-availability pair, which requires the profile to be the same on both units, you must:

1. Break high availability.
2. Assign the new profile to both units.
3. Re-establish high availability.

- Have the same number and types of interfaces.

  For the Firepower 4100/9300 chassis, all interfaces must be preconfigured in FXOS identically before you enable High availability. If you change the interfaces after you enable High availability, make the interface changes in FXOS on the Standby unit, and then make the same changes on the Active unit.

- Have these settings in a remote branch deployment:
  - Have the same data management interface to handle management traffic in a remote deployment.

    For example, if you have used eth0 in device 1, use the same interface (eth0) in device 2 as well.
  - Use data management interface for management traffic.

    You cannot have one unit managed using a data interface and the other using a management interface.

If you are using units with different flash memory sizes in your High availability configuration, make sure the unit with the smaller flash memory has enough space to accommodate the software image files and the configuration files. If it does not, configuration synchronization from the unit with the larger flash memory to the unit with the smaller flash memory will fail.

Software requirements

Software requirements are configuration specifications that High availability units must meet to function properly in a failover configuration.

Required software specifications

The two units in a High availability configuration must meet these requirements:

- Be in the same firewall mode (routed or transparent).
- Have the same software version.
- Be in the same group on the Cloud-Delivered Firewall Management Center.
- Have the same NTP configuration. See Synchronize time on FTD devices, on page 1303.
- Configuration changes must be fully deployed to the devices.

- Not have DHCP or PPPoE configured in any of their interfaces.
- (Firepower 4100/9300) Have the same flow offload mode, either both enabled or both disabled.

License requirements for Firewall Threat Defense devices in a high availability pair

Both Firewall Threat Defense units in a high availability configuration must have the same licenses.

High availability configurations require two license entitlements: one for each device in the pair.

Before high availability is established, it does not matter which licenses are assigned to the secondary/standby device. During high availability configuration, the Cloud-Delivered Firewall Management Center releases any unnecessary licenses assigned to the standby unit and replaces them with identical licenses assigned to the primary/active unit.

For example, if the active unit has a Essentials license and a IPS license, and the standby unit has only a Essentials license, the Cloud-Delivered Firewall Management Center communicates with the Smart Software Manager to obtain an available IPS license from your account for the standby unit. If your license account does not include enough purchased entitlements, your account becomes Out-of-Compliance until you purchase the correct number of licenses.

Failover and stateful failover links

A failover link and stateful failover link are dedicated connections that

- establish communication between the two units in a failover configuration
- require the same interface to be used between corresponding devices for consistency, and
- transmit critical failover and state information between the paired units.

Configuration recommendations

Cisco recommends to use the same interface between two devices in a failover link or a stateful failover link. For example, in a failover link, if you have used eth0 in device 1, use the same interface (eth0) in device 2 as well.

Failover link

A failover link is a communication mechanism that enables two units in a failover pair to constantly communicate to determine the operating status of each unit.

Failover link data

This below information is communicated over the failover link:

- The unit state (active or standby)
- Hello messages (keep-alives)
- Network link status
- MAC address exchange
- Configuration replication and synchronization

Interface for the failover link

A failover link interface is a dedicated communication channel that

- exists solely for failover communication between high availability units
- uses an unused data interface (physical, or EtherChannel)
- cannot be shared with normal networking interfaces or user data traffic, and
- requires specific interface sizing based on the device model.

Interface restrictions and requirements

You can use an unused data interface (physical, or EtherChannel) as the failover link; however, you cannot specify an interface that is currently configured with a name. You cannot use a data management interface if the interface is configured for communication with Security Cloud Control. You also cannot use a subinterface with the exception of a subinterface defined on the chassis for multi-instance mode. The failover link interface is not configured as a normal networking interface; it exists for failover communication only. This interface can only be used for the failover link (and also for the state link).

The Firewall Threat Defense does not support sharing interfaces between user data and the failover link. You also cannot use separate subinterfaces on the same parent for the failover link and for data (multi-instance chassis subinterfaces only). If you use a chassis subinterface for the failover link, then all subinterfaces on that parent, and the parent itself, are restricted for use as failover links.



Note When using an EtherChannel as the failover or state link, you must confirm that the same EtherChannel with the same member interfaces exists on both devices before establishing high availability.

See these guidelines for the failover link:

- Firepower 4100/9300—You cannot use the management-type interface for the failover link.
- See these guidelines for sizing the link.

Table 44: Failover link size

Model	Interface Size for Combined Failover and State Link
Firewall 200	1 Gbps
Firepower 1010	1 Gbps
Firepower 1100	1 Gbps
Secure Firewall 1200	1 Gbps
Secure Firewall 3100	Secure Firewall 3105—1 Gbps Secure Firewall 3110—1 Gbps Secure Firewall 3120—1 Gbps Secure Firewall 3130—10 Gbps Secure Firewall 3140—10 Gbps

Model	Interface Size for Combined Failover and State Link
Firepower 4100	10 Gbps
Secure Firewall 4200	10 Gbps
Firepower 9300	10 Gbps

The alternation frequency is equal to the unit hold time.



Note If you have a large configuration and a low unit hold time, alternating between the member interfaces can prevent the secondary unit from joining/re-joining. In this case, disable one of the member interfaces until after the secondary unit joins.

For an EtherChannel used as the failover link, to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link.

Connect the failover link

The failover link enables communication and coordination between units in a failover configuration.

You need to establish a physical connection between the failover units to enable proper failover functionality.

Procedure

Connect the failover link using one of the following methods:

- Using a switch, with no other device on the same network segment (broadcast domain or VLAN) as the failover interfaces of the Firewall Threat Defense device.
- Using an Ethernet cable to connect the units directly, without the need for an external switch.

If you do not use a switch between the units and the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which unit has the failed interface and caused the link to come down.

The failover link is physically connected and ready for failover configuration.

Stateful failover links

A stateful failover link is a network connection that

- passes connection state information between failover units, and
- must be configured to enable stateful failover functionality.

State link configuration requirements

To use Stateful Failover, you must configure a Stateful Failover link (also known as the state link) to pass connection state information.

Shared failover links

A shared failover link is an interface conservation method that enables multiple network functions to use the same failover connection, though dedicated interfaces for state and failover links should be considered for large configurations and high traffic networks.

Dedicated interface for the stateful failover link

A dedicated interface for the stateful failover link is a data interface configuration that uses a dedicated physical or EtherChannel interface specifically for the state link in failover deployments.

Performance considerations

You can use a dedicated data interface—such as a physical or EtherChannel interface—for the state link. For requirements about a dedicated state link, refer to Interface for the failover link, on page 685. See the section on connecting the state link for additional information Connect the failover link, on page 686.

For optimum performance with long-distance failover, the state link latency should be less than 10 milliseconds but no more than 250 milliseconds. If latency exceeds 10 milliseconds, performance might degrade due to retransmission of failover messages.

Failover and data link interruption avoidance

Failover and data link interruption avoidance is a network design practice that

- ensures failover links and data interfaces travel through different paths to decrease the chance that all interfaces fail at the same time
- prevents failover operation suspension when the failover link is down, and
- maintains network resilience by using proper connection scenarios.

Connection scenarios

We recommend that failover links and data interfaces travel through different paths to decrease the chance that all interfaces fail at the same time. If the failover link is down, the failover operation is suspended until the health of the failover link is restored.

These connection scenarios help design a resilient failover network.

Scenario 1—Not recommended

If a single switch or a set of switches are used to connect both failover and data interfaces between two Firewall Threat Defense devices, then when a switch or inter-switch-link is down, both Firewall Threat Defense devices become active. Therefore, the two connection methods shown in these figures are **not** recommended.

Figure 99: Connecting with a single switch—not recommended

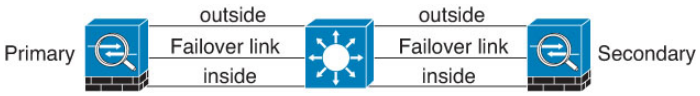
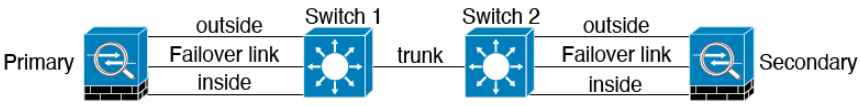


Figure 100: Connecting with a double-switch—not recommended



Scenario 2—Recommended

We recommend that failover links not use the same switch as the data interfaces. Instead, use a different switch or use a direct cable to connect the failover link, as shown in these figures.

Figure 101: Connecting with a different switch

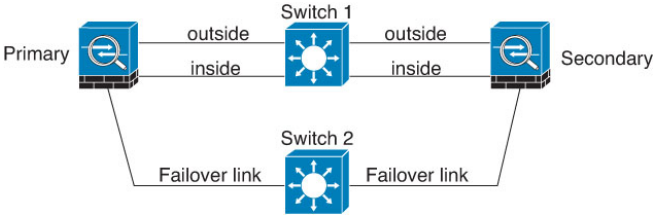
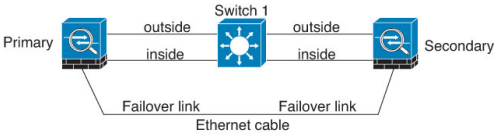


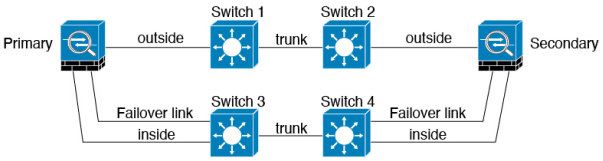
Figure 102: Connecting with a cable



Scenario 3—Recommended

If the Firewall Threat Defense data interfaces are connected to more than one set of switches, then a failover link can be connected to one of the switches, preferably the switch on the secure (inside) side of network, as shown in this figure.

Figure 103: Connecting with a secure switch



MAC addresses and IP addresses in High availability

MAC addresses and IP addresses in High availability are network addressing mechanisms that

- provide unique identification for network interfaces during failover scenarios
- maintain network connectivity when primary devices become unavailable, and
- ensure seamless traffic flow between active and standby configurations.

Address configuration types

When you configure your interfaces, you can specify an active IP address and a standby IP address on the same network. Generally, when a failover occurs, the new active unit takes over the active IP addresses and MAC addresses. Because network devices see no change in the MAC to IP address pairing, no ARP entries change or time out anywhere on the network.

Note Although recommended, the standby address is not required. Without a standby IP address, the active unit cannot perform network tests to check the standby interface health; it can only track the link state. You also cannot connect to the standby unit on that interface for management purposes.

The IP address and MAC address for the state link do not change at failover.

Active/standby configurations use specific address handling methods:

For Active/Standby High availability, see the following for IP address and MAC address usage during a failover event:

1. The active unit always uses the primary unit's IP addresses and MAC addresses.
2. When the active unit fails over, the standby unit assumes the IP addresses and MAC addresses of the failed unit and begins passing traffic.
3. When the failed unit comes back online, it is now in a standby state and takes over the standby IP addresses and MAC addresses.

However, if the secondary unit boots without detecting the primary unit, then the secondary unit becomes the active unit and uses its own MAC addresses, because it does not know the primary unit MAC addresses. When the primary unit becomes available, the secondary (active) unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic. Similarly, if you swap out the primary unit with new hardware, a new MAC address is used.

If you disable high availability and set the failover configurations to a disabled state, you will need to manually resume high availability, or reboot the device. It is recommended to use the command **configure high-availability resume** and resume the high availability instead of rebooting the device. If you reload the standby unit with the failover configuration disabled, the standby unit boots up as the active unit and uses the primary unit's IP addresses and MAC addresses. This leads to duplicate IP addresses and causes network traffic disruptions. Use the command **configure high-availability resume** to enable failover and restore the traffic flow.

Note If you enable failover on a standalone device, the data interfaces go down at negotiation state of failover, interrupting traffic.

Virtual MAC addresses guard against this disruption, because the active MAC addresses are known to the secondary unit at startup, and remain the same in the case of new primary unit hardware. We recommend that you configure the virtual MAC address on both the primary and secondary units to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not configure virtual MAC addresses, you might need to clear the ARP tables on connected routers to restore traffic flow. The Firewall Threat Defense does not send gratuitous ARPs for static NAT addresses

when the MAC address changes, so connected routers do not learn of the MAC address change for these addresses.

For Active/Active failover, see the following for IP address and MAC address usage during a failover event:

1. The primary unit autogenerates active and standby MAC addresses for all interfaces in failover group 1 and 2 contexts. You can also manually configure the MAC addresses if necessary, for example, if there are MAC address conflicts.
2. Each unit uses the active IP addresses and MAC addresses for its active failover group, and the standby addresses for its standby failover group. For example, the primary unit is active for failover group 1, so it uses the active addresses for contexts in failover group 1. It is standby for the contexts in failover group 2, where it uses the standby addresses.
3. When a unit fails over, the other unit assumes the active IP addresses and MAC addresses of the failed failover group and begins passing traffic.
4. When the failed unit comes back online, and you enabled the preempt option, it resumes the failover group.

Virtual MAC addresses provide specialized functionality:

The Firewall Threat Defense has multiple methods to configure virtual MAC addresses. We recommend using only one method. If you set the MAC address using multiple methods, the MAC address used depends on many variables, and might not be predictable.

For multi-instance capability, the FXOS chassis autogenerates only primary MAC addresses for all interfaces. You can overwrite the generated MAC address with a virtual MAC address with both the primary and secondary MAC addresses, but predefining the secondary MAC address is not essential; setting the secondary MAC address does ensure that to-the-box management traffic is not interrupted in the case of new secondary unit hardware.

MAC address table updates occur during failover events:

During failover, the device that is designated as the new active device generates multicast packets for each MAC address entry in the MAC table and sends them to all the bridge group interfaces. This action prompts the upstream switches in the bridge group to update their routing tables with the new active device's interface to ensure accurate traffic forwarding.

The time taken to generate multicast packets and update the routing tables of the upstream switches depends on the number of entries in the MAC address table and the number of bridge group interfaces. Use the **show failover statistics state-switch-delay** command to display statistics related to the delays encountered during failover events.

Stateful failover

A stateful failover is a high availability mechanism that

- continuously passes per-connection state information from the active unit to the standby unit
- maintains connection information after a failover occurs, and
- allows supported end-user applications to continue communication sessions without reconnecting.

Stateful failover operation

During Stateful Failover, the active unit continually passes per-connection state information to the standby unit. After a failover occurs, the same connection information is available at the new active unit. Supported end-user applications are not required to reconnect to keep the same communication session.

Supported features

Supported features are high availability capabilities that

- enable state information to be passed to the standby Firewall Threat Defense device during Stateful Failover
- maintain connection states and critical information for seamless traffic flow, and
- ensure minimal disruption during failover events.

State information types

For Stateful Failover, these state information types are passed to the standby device:

- NAT translation table.
- TCP and UDP connections and states, including HTTP connection states. Other types of IP protocols, and ICMP, are not parsed by the active unit, because they get established on the new active unit when a new packet arrives.
- Snort connection states, inspection results, and pin hole information, including strict TCP enforcement.
- The ARP table
- The Layer 2 bridge table (for bridge groups)
- The ISAKMP and IPsec SA table
- GTP PDP connection database
- SIP signaling sessions and pin holes.
- Static and dynamic routing tables—Stateful Failover participates in dynamic routing protocols, like OSPF and EIGRP, so routes that are learned through dynamic routing protocols on the active unit are maintained in a Routing Information Base (RIB) table on the standby unit. Upon a failover event, packets travel normally with minimal disruption to traffic because the active secondary unit initially has rules that mirror the primary unit. Immediately after failover, the re-convergence timer starts on the newly active unit. Then the epoch number for the RIB table increments. During re-convergence, OSPF and EIGRP routes become updated with a new epoch number. Once the timer is expired, stale route entries (determined by the epoch number) are removed from the table. The RIB then contains the newest routing protocol forwarding information on the newly active unit.



Note Routes are synchronized only for link-up or link-down events on an active unit. If the link goes up or down on the standby unit, dynamic routes sent from the active unit may be lost. This is normal, expected behavior.

- DHCP Server—DHCP address leases are not replicated. However, a DHCP server configured on an interface will send a ping to make sure an address is not being used before granting the address to a DHCP client, so there is no impact to the service. State information is not relevant for DHCP relay or DDNS.
- Access control policy decisions—Decisions related to traffic matching (including URL, URL category, geolocation, and so forth), intrusion detection, malware, and file type are preserved during failover. However, for connections being evaluated at the moment of failover, there are the following caveats:
  - AVC—App-ID verdicts are replicated, but not detection states. Proper synchronization occurs as long as the App-ID verdicts are complete and synchronized before failover occurs.
  - Intrusion detection state—Upon failover, once mid-flow pickup occurs, new inspections are completed, but old states are lost.
  - File malware blocking—The file disposition must become available before failover.
  - File type detection and blocking—The file type must be identified before failover. If failover occurs while the original active device is identifying the file, the file type is not synchronized. Even if your file policy blocks that file type, the new active device downloads the file.
- User identity decisions from the identity policy, including the user-to-IP address mappings gathered passively through ISE Session Directory, and active authentication through captive portal. Users who are actively authenticating at the moment of failover might be prompted to authenticate again.
- Network AMP—Cloud lookups are independent from each device, so failover does not affect this feature in general. Specifically:
  - Signature Lookup—If failover occurs in the middle of a file transmission, no file event is generated and no detection occurs.
  - File Storage—If failover occurs when the file is being stored, it is stored on the original active device. If the original active device went down while the file was being stored, the file does not get stored.
  - File Pre-classification (Local Analysis)—If failover occurs in the middle of pre-classification, detection fails.
  - File Dynamic Analysis (Connectivity to the cloud)—If failover occurs, the system might submit the file to the cloud.
  - Archive File Support—If failover occurs in the middle of an analysis, the system loses visibility into the file/archive.
  - Custom Blocking—If failover occurs, no events are generated.
- Security Intelligence decisions. However, DNS-based decisions that are in process at the moment of failover are not completed.
- RA VPN—Remote access VPN end users do not have to reauthenticate or reconnect the VPN session after a failover. However, applications operating over the VPN connection could lose packets during the failover process and not recover from the packet loss.
- From all the connections, only established ones will be replicated on the Standby device.

Unsupported features

Unsupported features are device capabilities that

- are not passed to the standby Firewall Threat Defense in Stateful Failover
- do not maintain state information during failover events, and
- may require re-establishment after a failover occurs.

State information limitations

For Stateful Failover, the following state information is not passed to the standby Firewall Threat Defense:

- Sessions in plaintext tunnels other than GREv0 and IPv4-in-IP. Sessions inside tunnels are not replicated and the new active node will not be able to reuse existing inspection verdicts to match the correct policy rules.
- Decrypted TLS/SSL connections—The decryption states are not synchronized, and if the active unit fails, then decrypted connections will be reset. New connections will need to be established to the new active unit. Connections that are not decrypted (in other words, those that match a TLS/SSL Do Not Decrypt rule action) are not affected and are replicated correctly.
- Multicast routing.

Bridge group requirements for high availability

Bridge group requirements for high availability are special considerations that

- address traffic loss on bridge group member interfaces during failover
- prevent switch ports from entering blocking state for 30 to 50 seconds, and
- provide workarounds to maintain connectivity during topology changes.

High availability workarounds

There are special considerations for high availability when using bridge groups.

When the active unit fails over to the standby unit, the switch port running Spanning Tree Protocol (STP) can go into a blocking state for 30 to 50 seconds when it senses the topology change. To avoid traffic loss on the bridge group member interfaces while the port is in a blocking state, you can configure one of the following workarounds:

- Switch port is in Access mode—Enable the STP PortFast feature on the switch:

  ```
  interface interface_id
    spanning-tree portfast
  ```

  The PortFast feature immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, it eventually transitions into STP blocking mode.

- If the switch port is in Trunk mode, or you cannot enable STP PortFast, then you can use one of the following less desirable workarounds that impacts failover functionality or STP stability:

- Disable interface monitoring on the bridge group and member interfaces.
- Increase the interface hold time in the failover criteria to a high value that will allow STP to converge before the unit fails over.
- Decrease the STP timers on the switch to allow STP to converge faster than the interface hold time.

Failover health monitoring

Failover health monitoring is a system capability that monitors each unit for overall health and interface health.

Health monitoring tests

The Firewall Threat Defense device performs tests to determine the state of each unit. This section includes information about how the Firewall Threat Defense device performs tests to determine the state of each unit.

Unit health monitoring

Unit health monitoring is a failover mechanism that

- determines the health of peer units by monitoring the failover link with hello messages
- sends LANTEST messages on each data interface when three consecutive hello messages are missed, and
- initiates appropriate failover actions based on peer unit responsiveness.

Unit health monitoring behavior

The Firewall Threat Defense device determines the health of the other unit by monitoring the failover link with hello messages. If a unit does not receive three consecutive hello messages on the failover link, the sends LANTEST messages on each data interface, including the failover link, to validate whether the peer is responsive. The action that the Firewall Threat Defense device takes depends on the response from the other unit. These are the possible actions:

- If the Firewall Threat Defense device receives a response on the failover link, then it does not fail over.
- If the Firewall Threat Defense device does not receive a response on the failover link, but it does receive a response on a data interface, then the unit does not failover. The failover link is marked as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby while the failover link is down.
- If the Firewall Threat Defense device does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.



Note During a high-availability failover event, the Firewall Threat Defense device device may briefly appear as **Offline** in the device's health monitoring dashboard. This happens because health alerts are cleared during the process and are only updated after the process is complete. Wait for the failover operation to finish.

Heartbeat module redundancy

Heartbeat module redundancy is a high availability feature that

- sends heartbeat messages using the data plane transport infrastructure to supplement control plane heartbeat packets
- increments a counter when the peer receives heartbeat packets in the data plane
- prevents false failover or split-brain scenarios when the control plane is congested with traffic.

Heartbeat module redundancy behavior

Each unit in the HA periodically sends a broadcast keepalive heartbeat packet over the failover link. If the control plane is too busy handling traffic, sometimes the heartbeat packets do not reach the peers, or the peers do not process the heartbeat packets due to CPU overloading. When peers cannot communicate the keepalive status within the configurable timeout period, a false failover or split-brain scenario occurs.

The heartbeat module in the data plane helps to avoid the occurrence of false failover or split-brain due to traffic congestion in the control plane.

- The additional heartbeat module works similarly to the control plane module but sends and receives heartbeat messages using the data plane transport infrastructure.
- When the peer receives heartbeat packets in the data plane, a counter gets incremented.
- If the heartbeat transfer in the control plane fails, the node checks the heartbeat counter in the data plane. If the counter is incrementing, then the peer is alive, and the cluster does not perform a failover in this situation.



Note

- The additional heartbeat module is enabled by default whenever HA is enabled. You do not have to set a poll interval for the additional heartbeat module in the data plane. This module uses the same heartbeat interval that you set for the control plane.

Interface monitoring

Interface monitoring is a failover capability that

- tracks the health of up to 1025 interfaces to detect interface failures
- performs interface tests when hello messages are not received within 15 seconds, and
- triggers failover when the number of failed interfaces meets the defined threshold.

Interface monitoring operation

When a unit does not receive hello messages on a monitored interface for 15 seconds, it runs interface tests. If one of the interface tests fails for an interface, but this same interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed, and the device stops running tests.

If the threshold you define for the number of failed interfaces is met (see **Devices** > **Device Management**, and then **High Availability** > **Failover Trigger Criteria**), and the active unit has more failed interfaces than

the standby unit, then a failover occurs. If an interface fails on both units, then both interfaces go into the "Unknown" state and do not count towards the failover limit defined by failover interface policy.

An interface becomes operational again if it receives any traffic. A failed device returns to standby mode if the interface failure threshold is no longer met.

If an interface has IPv4 and IPv6 addresses configured on it, the device uses the IPv4 addresses to perform the health monitoring. If an interface has only IPv6 addresses configured on it, then the device uses IPv6 neighbor discovery instead of ARP to perform the health monitoring tests. For the broadcast ping test, the device uses the IPv6 all nodes address (FE02::1).

Interface tests

Interface tests are monitoring mechanisms that:

- evaluate interface operational status during failover scenarios
- run sequentially with approximately 1.5 second duration per test, and
- determine interface failure status on Firewall Threat Defense device units.

Interface test sequence

The Firewall Threat Defense device uses these interface tests in sequence:

1. Link Up/Down test—A test of the interface status. If the Link Up/Down test indicates that the interface is down, then the device considers it failed, and testing stops. If the status is Up, then the device performs the Network Activity test.
2. Network Activity test—A received network activity test. At the start of the test, each unit clears its received packet count for its interfaces. As soon as a unit receives any eligible packets during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then the device starts the ARP test.
3. ARP test—A test for successful ARP replies. Each unit sends a single ARP request for the IP address in the most recent entry in its ARP table. If the unit receives an ARP reply or other network traffic during the test, then the interface is considered operational. If the unit does not receive an ARP reply, then the device sends a single ARP request for the IP address in the *next* entry in the ARP table. If the unit receives an ARP reply or other network traffic during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic, and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then the device starts the Broadcast Ping test.
4. Broadcast Ping test—A test for successful ping replies. Each unit sends a broadcast ping, and then counts all received packets. If the unit receives any packets during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic, and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then testing starts over again with the ARP test. If both units continue to receive no traffic from the ARP and Broadcast Ping tests, then these tests will continue running in perpetuity.

Interface status

Interface status is a monitoring indicator that

- shows the current operational state of monitored interfaces in a failover configuration
- tracks traffic reception and hello packet communication between peer units, and
- determines whether an interface is actively monitored by the failover process.

Interface status types

Monitored interfaces can have these status types:

- **Unknown**: Initial status. This status can also mean the status cannot be determined.
- **Normal**: The interface is receiving traffic.
- **Normal (Waiting)**: The interface is up, but has not yet received a hello packet from the corresponding interface on the peer unit.
- **Normal (Not-Monitored)**: The interface is up, but is not monitored by the failover process.
- **Testing**: Hello messages are not heard on the interface for five poll times.
- **Link Down**: The interface or VLAN is administratively down.
- **Link Down (Waiting)**: The interface or VLAN is administratively down and has not yet received a hello packet from the corresponding interface on the peer unit.
- **Link Down (Not-Monitored)**: The interface or VLAN is administratively down, but is not monitored by the failover process.
- **No Link**: The physical link for the interface is down.
- **No Link (Waiting)**: The physical link for the interface is down and has not yet received a hello packet from the corresponding interface on the peer unit.
- **No Link (Not-Monitored)**: The physical link for the interface is down, but is not monitored by the failover process.
- **Failed**: No traffic is received on the interface, yet traffic is heard on the peer interface.

Failover triggers and detection timing

Failover triggers and detection timing are high availability mechanisms that:

- initiate automatic switchover when specific failure conditions occur on the active unit,
- monitor system health through configurable detection thresholds and timing parameters, and
- ensure service continuity by transferring control to the standby unit when failures exceed defined limits.

Failover triggering events and timing parameters

These events trigger failover in a Firepower high availability pair:

- More than 50% of the Snort instances on the active unit are down.
- Disk space on the active unit is more than 90% full.

- The **no failover active** command is run on the active unit or the **failover active** command is run on the standby unit.
- The active unit has more failed interfaces than the standby unit.
- Interface failure on the active device exceeds the threshold configured.

  By default, failure of a single interface causes failover. You can change the default value by configuring a threshold for the number of interfaces or a percentage of monitored interfaces that must fail for the failover to occur. If the threshold is exceeded on the active device, failover occurs. If the threshold is exceeded on the standby device, the unit moves to **Fail** state.

  To change the default failover criteria, enter this command in global configuration mode:

Table 45: Interface policy command

Command	Purpose
failover interface-policy num [%] `hostname (config)# failover interface-policy 20%`	Changes the default failover criteria. When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

This table shows the failover triggering events and associated failure detection timing. If failover occurs, you can view the reason for the failover in the Message Center, along with various operations pertaining to the high availability pair. You can configure these thresholds to a value within the specified minimum-maximum range.

Table 46: Firewall Threat Defense failover times

Failover triggering event	Minimum	Default	Maximum
Active unit loses power, hardware goes down, or the software reloads or crashes. When any of these occur, the monitored interfaces or failover link do not receives any hello message.	800 milliseconds	15 seconds	45 seconds
Active unit interface physical link down.	500 milliseconds	5 seconds	15 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

Active/standby failover

Active/standby failover is a high availability configuration that

- lets you use a standby Firewall Threat Defense device to take over the functionality of a failed unit,
- enables the standby unit to become the active unit when the active unit fails, and
- allows the failed unit to come back online as the standby unit if the problem is temporary.

The failed unit must be replaced if the problem is not temporary. Set the standby unit to primary before replacing the failed unit to retain the configuration of the secondary unit.

For multiple context mode, the ASA can fail over the entire unit (including all contexts) but cannot fail over individual contexts separately.

Primary/secondary roles and active/standby status

Primary/secondary roles and active/standby status are failover configuration concepts that

- define unit hierarchy, with primary units taking precedence during simultaneous startup
- determine IP address and MAC address usage patterns during failover operations, and
- establish traffic flow responsibilities between paired units in Active/Standby failover configurations.

Primary and secondary unit differences

When setting up Active/Standby failover, you configure one unit to be primary and the other to be secondary. During configuration, the primary unit's policies are synchronized to the secondary unit. At this point, the two units act as a single device for device and policy configuration. However, for events, dashboards, reports and health monitoring, they continue to display as separate devices.

The main differences between the two units in a failover pair are related to which unit is active and which unit is standby, namely which IP addresses to use and which unit actively passes traffic.

However, a few differences exist between the units depending on which unit is primary (as specified in the configuration) and which unit is secondary:

- The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).
- The primary unit MAC addresses are always coupled with the active IP addresses. The exception to this rule occurs when the secondary unit becomes active and cannot obtain the primary unit MAC addresses over the failover link. In this case, the secondary unit MAC addresses are used.

Active unit determination at startup

Active unit determination at startup is a high availability process that

- sets a unit to standby if it boots and detects a peer already running as active,
- sets a unit to active if it boots and does not detect a peer, and
- sets the primary unit as active and secondary unit as standby if both units boot simultaneously.

Failover events

A failover event is a system condition that

- triggers an active unit to transfer control to a standby unit in Active/Standby failover configurations
- occurs on a unit basis, and
- follows specific failover policies that determine whether failover occurs and what actions each unit takes.

Failover event policies and actions

This table shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions.

Table 47: Failover events

Failure Event	Policy	Active Unit Action	Standby Unit Action	Notes
Active unit failed (power or hardware)	Failover	n/a	Become active Mark active as failed	No hello messages are received on any monitored interface or the failover link.
Formerly active unit recovers	No failover	Become standby	No action	None.
Standby unit failed (power or hardware)	No failover	Mark standby as failed	n/a	When the standby unit is marked as failed, then the active unit does not attempt to fail over, even if the interface failure threshold is surpassed.
Failover link failed during operation	No failover	Mark failover link as failed	Mark failover link as failed	You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.
Failover link failed at startup	No failover	Become active Mark failover link as failed	Become active Mark failover link as failed	If the failover link is down at startup, both units become active.
State link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Interface failure on active unit above threshold	Failover	Mark active as failed	Become active	None.
Interface failure on standby unit above threshold	No failover	No action	Mark standby as failed	When the standby unit is marked as failed, then the active unit does not attempt to fail over even if the interface failure threshold is surpassed.

Config-sync optimization

Config-sync optimization is a configuration synchronization enhancement that

- compares configuration hash values between active and joining devices to determine if full synchronization is necessary

- enables faster device rejoining by skipping full configuration synchronization when hash values match, and
- reduces maintenance window and upgrade time during high availability operations.

Guidelines and limitations of configuration sync optimization

When a device reboots or rejoins after a suspend or resume high availability, the joining device clears its running configuration. The active device then sends its entire configuration to the joining device for a full configuration synchronization. If the active device has a large configuration, this process can take several minutes.

The configuration sync optimization functionality enables comparing the configuration of the joining device and the active device by exchanging configuration hash values. If the hash computed on both active and joining devices match, the joining device skips full configuration synchronization and rejoin the high availability configuration. This functionality ensures faster peering and reduces maintenance window and upgrade time.

Guidelines and limitations:

- The configuration sync optimization functionality is enabled by default.
- Firewall Threat Defense multiple context mode supports configuration sync optimization by sharing the context order during full configuration synchronization, allowing comparison of context order during subsequent node-rejoin.
- If you configure passphrase and failover IPsec key, then configuration sync optimization is not effective as the hash value computed in the active and standby devices differs.
- If you configure the device with dynamic ACL or SNMPv3, configuration sync optimization is not effective.
- Active device synchronizes full configuration with flapping LAN links as default behavior. During failover flaps between active and standby devices, configuration sync optimization is not triggered and devices perform a full configuration synchronization.
- Configuration sync optimization gets triggered when the high availability configuration recovers from an interruption or loss of network communication between the active and standby devices.

Monitoring config-sync optimization:

- Syslog messages are generated displaying whether the hash values computed on the active and joining unit match, does not match, or if the operation timeout expires.
- The syslog message also displays the time elapsed, from the time of sending the hash request to the time of getting and comparing the hash response.



Note The Configuration State in **show failover state** command displays the config sync state status during HA joining. This state does not reflect the later config deployments or changes, and replication on device until a resync has been initiated.

Requirements and prerequisites for high availability

This reference provides the requirements and prerequisites needed to implement high availability functionality.

Model support

Secure Firewall Threat Defense

User roles

Admin

Guidelines for High availability

Model Support

- 200/ 1010/1210/ 1220:
  - You should not use the switch port functionality when using High availability. Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. High availability is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal High availability network setup, active switch ports on both units will lead to network loops. We suggest that you use external switches for any switching capability. Note that VLAN interfaces can be monitored by failover, while switch ports cannot. Though Cloud-Delivered Firewall Management Center allows you to configure switch ports / VLAN interfaces in HA, it may lead to misconfiguration due to network loops when connecting to external switches.
  - You can only use a firewall interface as the failover link.
- Firepower 9300—Intra-chassis High Availability is not supported.
- The Firewall Threat Defense Virtual on public cloud networks such as Microsoft Azure and Amazon Web Services are not supported with High availability because Layer 2 connectivity is required.

Additional Guidelines

- When the active unit fails over to the standby unit, the connected switch port running Spanning Tree Protocol (STP) can go into a blocking state for 30 to 50 seconds when it senses the topology change. To avoid traffic loss while the port is in a blocking state, you can enable the STP PortFast feature on the switch:

  interface *interface_id* **spanning-tree portfast**

  This workaround applies to switches connected to both routed mode and bridge group interfaces. The PortFast feature immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, the port eventually transitions into STP blocking mode.
- Configuring port security on the switches connected to the Firewall Threat Defense device failover pair can cause communication problems when a failover event occurs. This problem occurs when a secure

MAC address configured or learned on one secure port moves to another secure port, a violation is flagged by the switch port security feature.

- For Active/Standby High availability and a VPN IPsec tunnel, you cannot monitor both the active and standby units using SNMP over the VPN tunnel. The standby unit does not have an active VPN tunnel, and will drop traffic destined for the NMS. You can instead use SNMPv3 with encryption so the IPsec tunnel is not required.
- Both the peer devices go into unknown state and high-availability configuration fails if you run clish in any of the peer devices while creating a High Availability pair.
- Immediately after failover, the source address of syslog messages will be the failover interface address for a few seconds.
- For better convergence (during a failover), you must shut down the interfaces on a HA pair that are not associated with any configuration or instance.
- In high availability deployments with a shared failover and state interfaces, if the connectivity over failover interface fails while the standby unit is in sync config state, and eventually when the standby unit switches to active state, it results in a split brain. This is observed even when there is a monitored interface.

  The connectivity over failover interface can fail in the following situations:

  - After the standby unit reboots, the connectivity over failover interface fails while the standby unit is in sync config state.
  - After enabling previously disabled failover on the standby unit, the connectivity over failover interface fails while the standby unit is in sync config state.

  To avoid split brain, use failover and state interfaces. If the connectivity failures over the failover interfaces persist, isolate the new active (previously standby) unit from the network.
- If you configure failover encryption in evaluation mode, the systems use DES for the encryption. If you then register the devices using an export-compliant account, the devices will use AES after a reboot. Thus, if a system reboots for any reason, including after installing an upgrade, the peers will be unable to communicate and both units will become the active unit. We recommend that you do not configure encryption until after you register the devices. If you do configure this in evaluation mode, we recommend you remove the encryption before registering the devices.
- When using SNMPv3 with failover, if you replace a failover unit, then SNMPv3 users are not replicated to the new unit. You must remove the users, re-add them, and then redeploy your configuration to force the users to replicate to the new unit.
- The device does not share SNMP client engine data with its peer.
- If you have a very large number of access control and NAT rules, the size of the configuration can prevent efficient configuration replication, resulting in the standby unit taking an excessively long time to reach standby ready state. This can also impact your ability to connect to the standby unit during replication through the console or SSH session. To enhance configuration replication performance, enable transactional commit for both access rules and NAT, using the **asp rule-engine transactional-commit access-group** and **asp rule-engine transactional-commit nat** commands.
- A unit in a High availability pair transitioning to the standby role synchronizes its clock with the active unit.

  Example:

```
firepower#show clock
01:00:52 UTC Mar 1 2022

...
01:01:18 UTC Mar 1 2022 <======= Incorrect (previous) clock
Cold Standby               Sync Config               Detected an Active mate

19:38:21 UTC Apr 9 2022 <======= Updated clock
Sync Config                Sync File System          Detected an Active mate
...
firepower/sec/stby#show clock
19:38:40 UTC Apr 9 2022
```

- The units in High availability do not dynamically synchronize the clock. Here are some examples of events when synchronization takes place:
  - A new High availability pair is created.
  - High availability is broken and re-created.
  - Communication over the failover link was disrupted and reestablished.
  - Failover status was manually changed at the CLI using the **no failover/failover** or **configure high-availability suspend/resume** (Firewall Threat Defense) commands.
- Enabling High availability forces all routes to be deleted and are re-added after the High availability progression changes to the Active state. You could experience connection loss during this phase.
- If you replace the primary unit, then when you re-create high-availability, you should set the replacement unit as the *secondary* unit so that the configurations are replicated from the former secondary unit to the replacement unit. If you set the replacement unit as primary, you will accidentally overwrite the configuration that is present on the operational unit.
- Configure identical strong encryption on both Threat Defense devices before forming the High availability. If they are non-identical, the HA peers may enter a split-brain state.
- Deploying Firepower 1100 devices in high availability with hundreds of interfaces configured on them can result in increased delay in the failover time (seconds).
- In the High availability configuration, short-lived connections, usually using port 53, are closed quickly and never transferred or synchronized from Active to Standby, so there might be a difference in the number of connections on both High availability devices. This is expected behavior for short-lived connections. You can try to compare the connections that are long-lived (for example, more than 30-60 seconds).
- In the High availability configuration, embryonic connections—connection requests that have not yet completed the three-way handshake process—are closed quickly and not synchronized between the active and standby devices. This design ensures HA system efficiency and security. For this reason, there might be a difference in the number of connections on both High availability devices, which is to be expected.
- If the failover LAN link is not connected back-to-back and instead connected through one or more switches, a failure within the intermediate path can cause the active unit to lose connectivity with the standby unit, resulting in inconsistent active/standby states. Although this does not impact High availability functionality, it is recommended to check and recover the failover-link path between the active and standby units.

  When the failover LAN link is down, it is not recommended to deploy any configuration, as it may not be replicated to the peer unit.

- See the Secure Firewall Threat Defense Virtual getting started guides and review your Firewall Threat Defense Virtual device configurations for high availability.
- In OSPF, after a failover, the OSPF connection with the peer device becomes invalid. Terminate the invalid connection, and then establish a new OSPF connection.
- In transparent mode, if you have problems with the active unit losing the MAC address of the hot-standby router (HSRP), create a static mapping for the MAC address.
- UCAPL or CC compliance mode cannot be changed if the threat defense device is in high availability. Modify the compliance mode before forming the high availability pair.

Add a high availability pair

Establish an Active/Standby high-availability pair to ensure network protection continuity in case of device failure.

When establishing an Active/Standby high-availability pair, you designate one of the devices as primary and the other as secondary. The Cloud-Delivered Firewall Management Center deploys a merged configuration to the paired devices. If there is a conflict, the primary device setting is used.

Note The failover link and the stateful failover link are in a private IP space and are only used for communication between peers in a high-availability pair. After high availability is established, selected interface links and encryption settings cannot be modified without breaking the high-availability pair and reconfiguring it.

Caution Creating or breaking a high-availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196. The system warns you that continuing to create a high-availability pair restarts the Snort process on the primary and secondary devices and allows you to cancel.

Before you begin

Confirm that both devices:

- Are the same model.
- Have the same number and type of interfaces.
- Are in the same group.
- Have normal health status and are running the same software.
- Are either in routed or transparent mode.

Note Only routed mode is supported for manager access on a data interface.

- Have the same NTP configuration. See Synchronize time on FTD devices, on page 1303.
- Are fully deployed with no uncommitted changes.
- Do not have DHCP or PPPoE configured on any interfaces.
- For manager access on a data interface:
  - Use the same data interface on both devices for manager access.
  - Redundant manager access data interface is not supported.
  - You cannot use DHCP; only a static IP address is supported. Features that rely on DHCP cannot be used, including DDNS and zero-touch provisioning.

Note If you use zero-touch provisioning to register the device, when you use the outside interface for manager access, it uses DHCP by default. Before you can enable high availability, you need to change the IP address to a static address. See Change the device IP address, on page 1158. Alternatively, you can use the Management interface instead; DHCP is supported on Management with high availability.

- Have different static IP addresses in the same subnet.
- Use the same manager configuration (**configure manager add** command) to ensure that the connectivity is the same.
- You cannot use the data interface as the failover or state link.

Note The high availability formation is possible between the two Firewall Threat Defense devices when the certificate available on the primary device is not present on the secondary device. When high availability is formed, the certificate will be synched on the secondary device.

Procedure

Step 1 In the Security Cloud Control navigation bar, click **Security Devices**.

Step 2 Click the **Devices** tab to locate your device.

Step 3 Click the **FTD** tab and select the device you want to establish as the primary device.

Step 4 In the **Management** pane, click **High Availability**.

Step 5 Enter a display **Name** for the high-availability pair.

Step 6 Under **Device Type**, choose **Firepower Threat Defense**.

Step 7 Choose the **Primary Peer** device for the high-availability pair.

Step 8 Choose the **Secondary Peer** device for the high-availability pair.

Note

In the remote deployment, the devices appearing in the **Secondary Peer** list depend on the active device selected in the **Primary Peer** list:

- If the selected primary peer uses a data interface for management, only the data interface managed devices are listed in the secondary peer list.
- If the data management interface on the primary peer has an IPv4 address configured on it, then the secondary peer lists only the data interface managed devices that have an IPv4 address configured on them. The same rule applies to IPv6-managed devices as well.
- The data management interface names of primary and secondary devices should be the same. Devices with different interface names will not be listed in the secondary peer list.

Step 9 Click **Continue**.

Step 10 Under **LAN Failover Link**, choose an **Interface** with enough bandwidth to reserve for failover communications.

Note
Only interfaces that do not have a logical name, do not belong to a security zone, and are not used for handling management traffic, will be listed in the **Interface** drop-down in the **Add High Availability Pair** dialog.

Step 11 Type any identifying **Logical Name**.

Step 12 Type a **Primary IP** address for the failover link on the active unit.

This address should be on an unused subnet.

Note
169.254.1.0/24 and fd00:0:0:*::/64 are internally used subnets and cannot be used for the failover or state links.

Step 13 Optionally, choose **Use IPv6 Address**.

Step 14 Type a **Secondary IP** address for the failover link on the standby unit. This IP address must be in the same subnet as the primary IP address.

Step 15 If IPv4 addresses are used, type a **Subnet Mask** that applies to both the primary and secondary IP addresses.

Step 16 Optionally, under **Stateful Failover Link**, choose the same **Interface**, or choose a different interface and enter the high availability configuration information.

Note
169.254.1.0/24 and fd00:0:0:*::/64 are internally used subnets and cannot be used for the failover or state links.

Step 17 (Optional) Choose **Enabled** and choose the **Key Generation** method for IPsec Encryption between the failover links.

Step 18 Click **OK**. This process takes a few minutes as the process synchronizes system data.

After a successful configuration, the high availability pair is established with one device as active and the other as standby. You can see the **FTD High Availability** label on the Firewall Threat Defense node on the Security Cloud Control **Security Devices** page. Select the node to see the active and standby devices you configured for high availability.

Note We recommend to back up the devices. You can use the backup to quickly replace the devices when they fail and to restore the high availability service without being delinked from the Cloud-Delivered Firewall Management Center.

Configure optional high availability parameters

You can view the initial High Availability Configuration on the Cloud-Delivered Firewall Management Center. You cannot edit these settings without breaking the high availability pair and then re-establishing it. To edit these settings, you must first break the high availability pair and re-establish it. You can modify the Failover Trigger Criteria to improve failover results. Interface Monitoring helps you determine which interfaces are best suited for failover.

Configure standby IP addresses and interface monitoring

Configure standby IP addresses for high availability interfaces to enable the active unit to perform network tests and check standby interface health beyond just tracking link state.

For each interface, set a standby IP address. Although recommended, the standby address is not required. Without a standby IP address, the active unit cannot perform network tests to check the standby interface health; it can only track the link state.

By default, monitoring is enabled on all named physical interfaces and, for the 200/10101210/1220, all VLAN interfaces,. You might want to exclude interfaces attached to less critical networks from affecting your failover policy. Switch ports are not eligible for interface monitoring.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click the **Edit** ( ).

Step 3 Click the **High Availability** tab.

Step 4 In the **Monitored Interfaces** area, click the **Edit** ( ) next to the interface you want to edit.

Step 5 Check the **Monitor this interface for failures** check box.

Step 6 On the **IPv4** tab, enter the **Standby IP Address**.

This address must be a free address on the same network as the active IP address.

Step 7 If you configured the IPv6 address manually, on the **IPv6** tab, click the **Edit** ( ) next to the active IP address, enter the **Standby IP Address**, and click **OK**.

This address must be a free address on the same network as the active IP address. For autogenerated and **Enforce EUI 64** addresses, the standby address is automatically generated.

Step 8 Click **OK**.

Edit high availability failover criteria

This task allows you to customize failover criteria for high availability device pairs to match your specific network deployment requirements.

You can customize failover criteria based on your network deployment.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click the **Edit** ( ).

Step 3 Choose **High Availability**.

Step 4 Next to **Failover Trigger Criteria**, click the **Edit** ( ).

Step 5 Under **Interface Failure Threshold**, choose the number or percentage of interfaces that must fail before the device fails over.

Step 6 Under **Hello packet Intervals**, choose how often hello packets are sent over the failover link.

Step 7 Click **OK**.

Configure virtual MAC addresses

Configure active and standby MAC addresses to minimize loss of traffic during failover and provide redundancy against IP address mapping for failover.

You can configure active and standby MAC addresses for failover using the following methods in the Cloud-Delivered Firewall Management Center:

- From the **Advanced tab** on the **Edit Interface** page during interface configuration; see Configure the MAC address, on page 1110.
- From the **Add Interface MAC Address** dialog-box which is accessed from the **High Availability** page; see this procedure.

Note To configure the MAC address in both primary and secondary units (so that the MAC address is transferred to all sub-interfaces to both the high-availability units), the recommended approach is to use the **Interfaces** tab to replicate the MAC addresses on sub-interfaces over both active and standby high-availability units.

If you configure active and standby MAC addresses in both locations, the addresses defined during interface configuration take precedence for failover.

You can minimize loss of traffic during failover by designating active and standby MAC addresses to the physical interface. This feature offers redundancy against IP address mapping for failover.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

Step 3 Click **High Availability**.

Step 4 Click the **Add** ( ) icon next to Interface MAC Addresses.

Step 5 Choose a **Physical Interface**.

Step 6 Enter the **Active Interface MAC Address**.

Step 7 Enter the **Standby Interface MAC Address**.

Step 8 Click **OK**.

Note
For detailed information, see Task 2, steps from 10 to 14 in Configure FTD High Availability on Firepower Appliances.

Manage High availability

This section describes how to manage High availability units after you enable High availability, including how to change the High availability setup and how to force failover from one unit to another.

Switch the Active Peer in the Firewall Threat Defense High Availability Pair

This task allows you to manually switch the active and standby units in a high availability pair, effectively forcing failover for reasons such as persistent fault or health events on the current active unit.

After you establish the Firewall Threat Defense high availability pair, you can manually switch the active and standby units, effectively forcing failover for reasons such as persistent fault or health events on the current active unit. Both units should be fully deployed before you complete this procedure.

Before you begin

Refresh node status for a single Firewall Threat Defense high availability pair, on page 711. This ensures that the status on the Firewall Threat Defense high availability device pair is in sync with the status on the Cloud-Delivered Firewall Management Center.

Follow these steps to switch the active peer in the high availability pair:

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high availability pair where you want to change the active peer, click the **Switch Active Peer**.

Step 3 Choose an option:

- Click **Yes** to immediately make the standby device the active device in the high availability pair.
- Click **No** to cancel and return to the Device Management page.

Refresh node status for a single Firewall Threat Defense high availability pair

Refresh the high availability node status to obtain accurate information about the active and standby device in a high availability pair when communication issues result in out of sync data.

Whenever active or standby devices in the Firewall Threat Defense high availability pair are rebooted, the Cloud-Delivered Firewall Management Center may not display accurate high availability status for either device. This is because when the device reboots, the high availability status is immediately updated on the device and its corresponding event is sent to the Cloud-Delivered Firewall Management Center. However, the status may not be updated on the Cloud-Delivered Firewall Management Center because the communication between the device and the Cloud-Delivered Firewall Management Center is yet to be established.

Communication failures or weak communication channels between the Cloud-Delivered Firewall Management Center and devices may result in out of sync data. When you switch the active and standby devices in a high availability pair, the change may not be reflected in the Cloud-Delivered Firewall Management Center even after a significant time duration.

In these scenarios, you can refresh the high availability node status to obtain accurate information about the active and standby device in a high availability pair.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high availability pair where you want to refresh the node status, click the **Refresh HA Node Status**.

Step 3 Click **Yes** to refresh the node status.

Suspend and resume high availability

You can suspend a unit in a high-availability pair, which is useful when:

- Both units are in an active-active situation, and fixing the communication on the failover link does not correct the problem.
- You want to troubleshoot an active or standby unit and do not want the units to fail over during that time.

When you suspend high availability, the currently active device remains active and continues handling all user connections. However, failover criteria are no longer monitored, and the system will never fail over to the now pseudo-standby device.

When using a data interface for manager access, the management connection will be broken until after you resume.

The primary difference between suspending high availability and breaking high availability is that on a suspended high-availability device, the high-availability configuration is retained. When you break high availability, the configuration is erased. Thus, you have the option to resume high availability on a suspended system, which enables the existing configuration and makes the two devices function as a failover pair again.

To suspend high availability, use the **configure high-availability suspend** command.

```
> configure high-availability suspend
Please ensure that no deployment operation is in progress before suspending
high-availability.
Please enter 'YES' to continue if there is no deployment operation in
progress and 'NO' if you wish to abort: YES
Successfully suspended high-availability.
```

If you suspend high availability from the active unit, the configuration is suspended on both the active and standby units. The standby unit interface configuration is also erased. If you suspend it from the standby unit, it is suspended on the standby unit only, but the active unit will not attempt to fail over to a suspended unit.

To resume failover, use the **configure high-availability resume** command.

```
> configure high-availability resume
Successfully resumed high-availablity.
```

You can resume a unit only if it is in Suspended state. The unit will negotiate active/standby status with the peer unit.

Note Suspending high availability is a temporary state. If you reload a unit, it resumes the high-availability configuration automatically and negotiates the active/standby state with the peer.

Replace a unit in Firewall Threat Defense high availability pair

This reference explains the process for replacing a failed unit in a Firewall Threat Defense high availability pair, including both backup restore and manual replacement procedures.

To replace a failed unit in the Firewall Threat Defense high availability pair using a backup file, see *Restoring Cloud-Delivered Firewall Management Centers and Managed Devices* in the *Cisco Secure Firewall Management Center Administration Guide*.

If you do not have a backup of the failed device, you must break high availability. Then, register the replacement device to the Cloud-Delivered Firewall Management Center and reestablish high availability. The process varies depending on whether the device is primary or secondary:

- Replace a Primary Firewall Threat Defense HA Unit with no Backup, on page 712
- Replace a Secondary Firewall Threat Defense HA Unit with no Backup, on page 713

Replace a Primary Firewall Threat Defense HA Unit with no Backup

This task replaces a failed primary unit in a Firewall Threat Defense high availability pair to restore redundancy and maintain network protection.

Follow these steps below to replace a failed primary unit in the Firewall Threat Defense high availability pair. Failing to follow these steps can overwrite the existing high availability configuration.

Caution Creating or breaking the Firewall Threat Defense high availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196. The system warns you that continuing to create a high availability pair restarts the Snort process on the primary and secondary devices and allows you to cancel.

Caution Never move a disk from sensor or Cloud-Delivered Firewall Management Center to another device without reimaging the disk. This is an unsupported configuration and can cause breakage in functionality.

Procedure

Step 1 Choose **Force Break** to separate the high availability pair; see Break a high availability pair, on page 714.

Note
The break operation removes all the configuration related to HA from Firewall Threat Defense and Cloud-Delivered Firewall Management Center, and you need to recreate it manually later. To successfully configure the same HA pair, ensure that you save the IPs, MAC addresses, and monitoring configuration of all the interfaces/subinterfaces prior to executing the HA break operation.

Step 2 Unregister the failed primary Firewall Threat Defense device from the Cloud-Delivered Firewall Management Center.

Step 3 Register the replacement Firewall Threat Defense to the Cloud-Delivered Firewall Management CenterPrerequisites to Onboard a Device to Cloud-Delivered Firewall Management Center, on page 21.

Step 4 Configure high availability, using the existing secondary/active unit as the primary device and the replacement device as the secondary/standby device during registration; see Add a high availability pair, on page 705.

Replace a Secondary Firewall Threat Defense HA Unit with no Backup

This task allows you to restore high availability functionality by replacing a failed secondary unit in an Firewall Threat Defense high availability pair.

When a secondary unit fails in a high availability configuration, you need to replace it and recreate the HA pair to maintain network redundancy and protection.

Caution Creating or breaking the Firewall Threat Defense high availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196. The system warns you that continuing to create a high availability pair restarts the Snort process on the primary and secondary devices and allows you to cancel.

Procedure

Step 1 Choose **Force Break** to separate the high availability pair.

See Break a high availability pair, on page 714.

Note
The break operation removes all the configuration related to HA from Firewall Threat Defense and Cloud-Delivered Firewall Management Center, and you need to recreate it manually later. To successfully configure the same HA pair, ensure that you save the IPs, MAC addresses, and monitoring configuration of all the interfaces/subinterfaces prior to executing the HA break operation.

Step 2 Unregister the secondary Firewall Threat Defense device from the Cloud-Delivered Firewall Management Center

Step 3 Register the replacement Firewall Threat Defense to the Cloud-Delivered Firewall Management Center.

See Prerequisites to Onboard a Device to Cloud-Delivered Firewall Management Center, on page 21.

Step 4 Configure high availability, using the existing primary/active unit as the primary device and the replacement device as the secondary/standby device during registration.

See Add a high availability pair, on page 705.

Break a high availability pair

Break a high availability pair to remove the high-availability configuration from both units while maintaining traffic flow on the active unit.

When you break a high-availability pair, the high-availability configuration is removed from both units.

When using the Management interface for manager access: The active unit remains up and passing traffic. The standby unit interface configuration is erased.

When using a data interface for manager access: See the following details.

- The active unit remains up and passing traffic.
- The standby unit data interfaces are shut down except for the manager access interface, which remains up using the standby IP address so it can maintain the management connection.

- In a remote branch deployment setup, all standby unit data interfaces that are assigned with a logical name are shut down except for the manager access interface, which remains up to maintain the management connection.
- If the primary unit is in the standby state:
  - The IP addresses for manager access are swapped permanently in the Cloud-Delivered Firewall Management Center configuration: the primary unit uses the standby IP address, and the secondary unit uses the active IP address.

Policies that were not deployed to the active unit prior to the break operation continue to remain un-deployed after the break operation is complete. Deploy the policies on the standalone device after the break operation is complete.

Note

- When IPsec is enabled on high availability interfaces on the Firewall Threat Defense device, the device cannot prioritize the encrypted packets into the high-priority receive queue. As a result, during high-volume data traffic scenarios, attempts to break high availability may fail as the device cannot efficiently manage and prioritize the large number of encrypted connections. To view the device's resource usage and the maximum throughput, use the `show resource usage` command.
- If you cannot reach the high-availability pair using the Cloud-Delivered Firewall Management Center, connect to the CLI on each device and enter **configure high-availability disable** to manually break high availability. See also Remove a High Availability Pair , on page 716.

Caution Breaking the Firewall Threat Defense high-availability pair immediately restarts the Snort process on the primary and secondary units, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Before you begin

- Refresh node status for a single Firewall Threat Defense high availability pair, on page 711. This ensures that the status on the high-availability pair is in sync with the status on the Cloud-Delivered Firewall Management Center.

Follow these steps to break a high availability pair:

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 If you have switched to the new device management page, select the node to open the side pane, and under **HA actions**, click **Break**.

Step 3 Next to the high-availability pair you want to break, click **More** (⋮) and choose **Break**.

Step 4 If the standby peer does not respond, check **Force Break**.

Step 5 Click **Yes**.

The Break operation removes the high-availability configuration from the active and standby units.

A FlexConfig policy deployed on the active unit may show a deployment failure after the break high-availability operation. You must alter and re-deploy the FlexConfig policy on the active unit.

What to do next

If you are using a FlexConfig policy on the active unit, alter and re-deploy the FlexConfig policy to eliminate deployment errors.

Note After you break high availability, the Firewall Threat Defense device that was operating as the active unit will still have the standby unit's IP address listed in its configuration. To resolve this, do an additional deployment on the formerly active Firewall Threat Defense device to remove the standby unit's IP address from its configuration.

Remove a High Availability Pair

This task allows you to unregister a High Availability pair from the Cloud-Delivered Firewall Management Center, which keeps the High Availability pair intact. You might want to unregister the pair if you want to register it to a new Cloud-Delivered Firewall Management Center or if the Cloud-Delivered Firewall Management Center can no longer reach the pair.

Unregistering a High Availability pair:

- Severs all communication between the Cloud-Delivered Firewall Management Center and the pair.
- Removes the pair from the **Device Management** page.
- Returns the pair to local time management if the pair's platform settings policy is configured to receive time from the Cloud-Delivered Firewall Management Center using NTP.
- Leaves the configuration intact, so the pair continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the pair again to the same or a different Cloud-Delivered Firewall Management Center causes the configuration to be removed, so the pair will stop processing traffic at that point; the High Availability configuration remains intact so you can add the pair as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.

Before you begin

- This procedure requires CLI access to the primary unit.

Procedure

Step 1 Log into Security Cloud Control and click **Security Devices**.

Step 2 Click the **FTD** tab and locate the High Availability pair you want to unregister. Select it so the device row is highlighted.

Step 3 In the **Device Actions** pane located to the right, click **Remove**.

Step 4 When prompted, select **OK** to confirm the removal of the selected device.

Monitoring High availability

Monitoring High availability is a system capability that enables you to view the High availability status.

View failover history

View the failover history of both high availability devices in a single view to understand the chronological order of failover events and their reasons.

You can view the failover history of both high availability devices in a single view. The history displays in chronological order and includes the reason for any failover.

Before you begin

Follow these steps to view failover history:

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

Step 3 Choose **Summary**.

Step 4 Under General, click **View** ( ).

View stateful failover statistics

This task allows you to monitor the stateful failover link statistics to assess the performance and health of your high availability configuration.

You can view the stateful failover link statistics of both the primary and secondary devices in the high availability pair.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

Step 3 Choose **High Availability**.

Step 4 Under Stateful Failover Link, click **View** ( ).

Step 5 Choose a device to view statistics.

High availability break troubleshooting in remote branch deployments

High availability break troubleshooting in remote branch deployments is a process that

- addresses common problems when breaking a High Availability pair in a remote deployment
- resolves active-active state issues, and
- fixes connectivity and failover link problems between devices and Security Cloud Control.

Common troubleshooting scenarios

These scenarios describe common problems you may encounter when breaking a High Availability pair in a remote deployment:

- Both Units are in Active-Active State.
- The primary or secondary device has lost connectivity with Security Cloud Control, and the failover link has become non-operational.
- The secondary device is in failed or disabled state and has lost connectivity with Security Cloud Control.

Break a high availability pair in active-active state

This task resolves the unstable network condition that occurs when both units in a high availability pair use the active IP address on their data management interface, resulting in network instability between the units and Security Cloud Control.

Both units in a remote deployment are in an active-active state because the failover interface became non-operational and they stopped receiving a response on their data interfaces. In this case, both units use the active IP address on their data management interface, which results in an unstable network between the units and Security Cloud Control.

You can determine if the units are both in active mode by logging into the device CLI and using the "show failover state" command on both units. The device status of both units shows 'active', and the same active IP address is assigned to both units.

Note You can try rectifying the failover interface to restore the communication between the two peers and then perform the **Force Break** operation.

If you cannot repair the connectivity issues of the failover interface, then perform these steps.

Procedure

Step 1 Identify a device you want to remove from the network among the two units.

Step 2 Connect to the CLI of the identified device, either from the console port or using SSH.

Step 3 Log in with the Admin username and password.

Step 4 Enter the **pmtool disablebyid sftunnel** command.

Note
Only use **pmtool** commands under the direction of the Cisco Technical Assistance Center.

Step 5 Disconnect all the interfaces from the device you want to remove from the network.

Step 6 Enter **configure network management-data-interface ipv4 manual** *ip_address ipv4_netmask gateway_ip_address* **interface***interface_id* command.

In *ip_address*specify the IP address of the standby device.

Example:

```
Configure network management-data-interface ipv4 manual 10.10.6.7 255.255.255.0 interface
gig0/0
Configuration updated successfully..!!
```

Step 7 Enter **configure high-availability suspend** to suspend HA.

```
configure high-availability suspend
Please ensure that no deployment operation is in progress before suspending
high-availability.
Please enter 'YES' to continue if there is no deployment operation in
progress and 'NO' if you wish to abort: YES
Successfully suspended high-availability.
```

Step 8 In the Security Cloud Control navigation bar, click **Security Devices**.

Step 9 Click the **Devices** tab to locate your device.

Step 10 Click the **FTD** tab and select the primary device.

Step 11 In the **Management** pane on the left, click **High Availability**.

Step 12 Choose **Device** > **Device Management**.

Step 13 Next to the high availability pair where you want to separate the high availability pair, click **Force Break**.

A message is displayed that the high-availability pair is separated successfully.

Step 14 Connect all the interfaces to the device.

Step 15 At the FTD CLI, enter **pmtool enablebyId sftunnel**.

The threat defense device establishes its connection with Security Cloud Control in sometime.

Note
It may take up to 5 minutes for the device to establish communication with Security Cloud Control.

Step 16 Enter the **sftunnel-status-brief** command to view the management connection status.

```
sftunnel-status-brief
PEER:10.10.17.202
Registration: Completed.
Connection to peer '10.10.17.202' Attempted at Wed Feb 9 09:21:57 2020 UTC
Last disconnect time : Wed Feb 9 09:19:09 2020 UTC
```

Step 17 Choose **Deploy** > **Deployment** to deploy the changes.

Before the Security Cloud Control deploys the changes, it will detect the configuration differences and stop the deployment. Security Cloud Control detects the IP address change made to the device outside of the Security Cloud Control.

Step 18 Synchronize interface changes with Security Cloud Control. See Sync Interface Changes with the Cloud-Delivered Firewall Management Center, on page 1014.

Step 19 You can now deploy the pending changes to the device. See Deploy configuration changes..

What to do next

(optional) Deploy any pending changes to the other device having the IP address of the active device.

Break a high availability pair when active or standby unit has lost connectivity

This task breaks a high availability pair when communication issues prevent normal failover operations between the active and standby units.

Problem: One of the peers has lost connectivity with Cloud-Delivered Firewall Management Center, and the failover link has become non-operational.

Table 48: Scenario:

Primary Device State	Secondary Device Stat	Primary Device Connectivity with Security Cloud Control?	Secondary Device Connectivity with ?Security Cloud Control	Failover link Operational? (Connectivity between Primary and Secondary devices)
Active	Standby	Yes	No	No
Standby	Active	No	Yes	No

Solution:

First, you can try rectifying the failover interface to restore the communication between the two peers and then perform the break or force break operation to separate the units.

If you cannot repair the connectivity issues of the failover interface, then you must complete additional steps using the device CLI after performing a high availability break operation.

Procedure

Step 1 In the Security Cloud Control navigation bar, click **Security Devices**.

Step 2 Click the **Devices** tab to locate your device.

Step 3 Click the **FTD** tab and select the primary device.

Step 4 In the **Management** pane on the left, click **High Availability**.

Step 5 Choose **Devices** > **Device Management**.

Step 6 Next to the high-availability pair you want to break, click the **Break HA**.

Step 7 Optionally, you can also check the check box to force break as one of the peers does not respond.

Step 8 Click **Yes**.

Step 9 Delete the standby device from Security Cloud Control.

a) Choose **Devices** > **Device Management**.

b) Next to the device you want to delete, click **Delete**.

What to do next

You can onboard the device as a standalone device to Security Cloud Control if required.

Break a high availability pair when the secondary device is in a failed or disabled state

This task separates high availability devices when the secondary device cannot communicate with Security Cloud Control and allows the primary device to operate independently.

Use this procedure when the secondary device is in a failed or disabled state and has lost connectivity with Security Cloud Control. In addition, the failover link may or may not be operational.

Table 49: Scenario:

Primary Device State	Secondary Device Stat	Primary Device Connectivity with Security Cloud Control?	Secondary Device Connectivity with Security Cloud Control?	Failover link Operational? (Connectivity between Primary and Secondary devices)
Active	Failed	Yes	No	Yes or No
Active	Disabled	Yes	No	Yes or No

Perform a high availability force break to separate the units and then use the device CLI to remove the configuration from the standby unit and make the device a standalone device.

Procedure

Step 1 In the Security Cloud Control navigation bar, click **Security Devices**.

Step 2 Click the **Devices** tab to locate your device.

Step 3 Click the **FTD** tab and select the primary device.

Step 4 In the **Management** pane on the left, click **High Availability**.

Step 5 Choose **Devices** > **Device Management**.

Step 6 Next to the high-availability pair you want to break, click the **Break HA**.

Step 7 Check the check box to force break as one of the peers does not respond.

Step 8 Click **Yes**.

Step 9 Delete the standby device from Security Cloud Control.

a) Choose **Devices** > **Device Management**.
b) Next to the device you want to delete, click **Delete**.

Step 10 Connect to the standby device's CLI, either from the console port or using SSH.

Step 11 Log in with the Admin username and password.

Step 12 Enter **configure high-availability disable** to remove the failover configuration and disable the data management interface on the device.

Step 13 Enter **configure network management-data-interface**.

Example:

```
configure network management-data-interface
Data interface to use for management: ethernet1/1
Specify a name for the interface [outside]: internet
IP address (manual / dhcp) [dhcp]: manual
IPv4/IPv6 address: 10.10.6.7
Netmask/IPv6 Prefix: 255.255.255.0
Default Gateway: 10.10.6.1
Comma-separated list of DNS servers [none]: 208.67.222.222,208.67.220.220
DDNS server update URL [none]:
Do you wish to clear all the device configuration before applying ? (y/n) [n]:

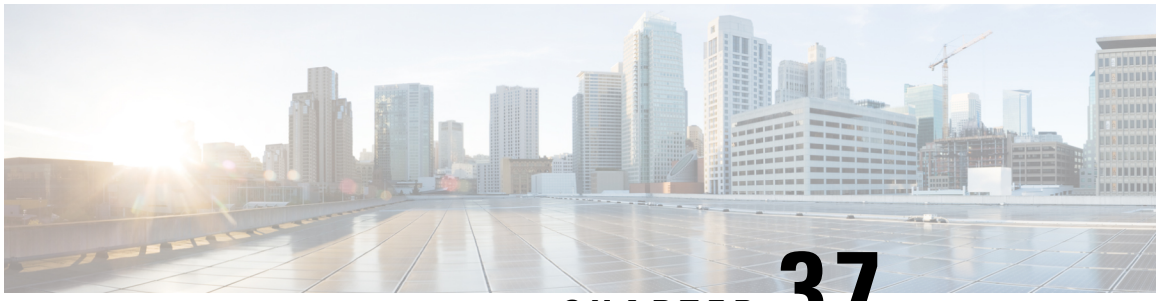
Configuration done with option to allow FMC access from any network, if you wish to change
 the FMC access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.
```

The new network settings are assigned to the data device.

What to do next

You can onboard the device as a standalone device to Security Cloud Control if required.



CHAPTER 37

Clustering: Secure Firewall 3100/4200

Clustering lets you group multiple Firewall Threat Defense nodes together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices.

Note Some features are not supported when using clustering. See Unsupported features with clustering, on page 782.

- Clustering for the Secure Firewall 3100/4200, on page 723
- Licenses for clustering, on page 725
- Prerequisites for clustering, on page 725
- Guidelines for clustering, on page 726
- Clustering configuration, on page 731
- Cluster node management, on page 762
- Monitor the cluster, on page 772
- Troubleshoot the cluster, on page 777
- Clustering examples, on page 780
- Reference for Clustering, on page 782
- History for Clustering, on page 796

Clustering for the Secure Firewall 3100/4200

This section describes the clustering architecture and how it works.

How a cluster fits into your network

The cluster consists of multiple firewalls acting as a single unit. To act as a cluster, the firewalls need the following infrastructure:

- Isolated, high-speed backplane network for intra-cluster communication, known as the *cluster control link*.
- Management access to each firewall for configuration and monitoring.

Infrastructure requirements

When you place the cluster in your network, the upstream and downstream routers need to be able to load-balance the data coming to and from the cluster using one of these methods:

- **Spanned EtherChannel** (Recommended)—Interfaces on multiple members of the cluster are grouped into a single EtherChannel; the EtherChannel performs load balancing between units.
- **Policy-Based Routing** (Routed firewall mode only)—The upstream and downstream routers perform load balancing between units using route maps and ACLs.
- **Equal-Cost Multi-Path Routing** (Routed firewall mode only)—The upstream and downstream routers perform load balancing between units using equal cost static or dynamic routes.

Control and data node roles

One member of the cluster is the control node. If multiple cluster nodes come online at the same time, the control node is determined by the priority setting; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. When you first create the cluster, you specify which node you want to be the control node, and it will become the control node simply because it is the first node added to the cluster.

All nodes in the cluster share the same configuration. The node that you initially specify as the control node will overwrite the configuration on the data nodes when they join the cluster, so you only need to perform initial configuration on the control node before you form the cluster.

Some features do not scale in a cluster, and the control node handles all traffic for those features.

Cluster interfaces

You can configure data interfaces as either Spanned EtherChannels or as Individual interfaces. All data interfaces in the cluster must be one type only. Refer Cluster interfaces, on page 731 for more information.

Interface type restrictions

- For Spanned EtherChannels: You can use regular firewall interfaces or IPS-only interfaces (inline sets or passive interfaces).
- For Individual interfaces: IPS-only interfaces are not supported.

Cluster control links

A cluster control link is a dedicated hardware interface that each unit must allocate for cluster communication. For more information, refer Cluster control link, on page 731.

Configuration replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Management network

You must manage each node using the Management interface; management from a data interface is not supported with clustering.

Licenses for clustering

Lists the licensing requirements for clustering.

Feature license assignment

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the Cloud-Delivered Firewall Management Center, you can specify the feature licenses you want to use for the cluster. Before you create the cluster, it doesn't matter which licenses are assigned to the data nodes; the license settings for the control node are replicated to each of the data nodes. You can modify licenses for the cluster by clicking **Edit Licenses** in **System** (⚙) > **Licenses** > **Smart Licenses** or choosing **Devices** > **Device Management**, clicking **Edit** (✎) for the cluster, and then in the **License** area, clicking **Edit** (✎).

APEX, PLUS, and VPN-Only licenses are not supported for clustering. Before you create a cluster, remove unsupported license assignments from the devices that will become cluster nodes.

Note If you add the cluster before the Cloud-Delivered Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Cloud-Delivered Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Prerequisites for clustering

Review these prerequisites before configuring clustering on Secure Firewall devices.

Model requirements

- Secure Firewall 3100—Maximum 16 nodes
- Secure Firewall 4200—Maximum 16 nodes

User roles

- Admin
- Access Admin
- Network Admin

Hardware and software requirements

- All units in a cluster must be the same model.
- All units must include the same number of interfaces and the same network modules, and must use the same network module slot assignments.
- The Cloud-Delivered Firewall Management Center access must be from the Management interface; data interface management is not supported.
- All units must run the identical software except at the time of an image upgrade. Hitless upgrade is supported.
- All units must be in the same firewall mode, routed or transparent.
- All units must use the same compliance mode.
- If core allocation is supported, all units must use the same core allocation profile.
- All units must be in the same domain and the same group, with no deployment pending or in progress.
- The control node must not have any unsupported features configured (refer Unsupported features with clustering, on page 782), and data nodes must not have any VPN configured. The control node can have site-to-site VPN configured.



Note

- Clustering is not supported for Threat Defense instances, such as instances in container mode.
- All units must use the same compliance mode.
- If core allocation is supported, all units must use the same core allocation profile.

Switch requirements

Be sure to complete the switch configuration before you configure clustering. Make sure the ports connected to the cluster control link have the correct (higher) MTU configured. By default, the cluster control link MTU is set to 100 bytes higher than the data interfaces. If the switches have an MTU mismatch, the cluster formation will fail.

Guidelines for clustering

Configuration requirements

Ensure that firewall mode configuration matches on all units in the cluster.

- High Availability is not supported with clustering.
- The cluster control link is only supported using IPv4.

Switch and EtherChannel configuration

Disable the LACP Graceful Convergence feature on all cluster-facing EtherChannel interfaces for Cisco Nexus switches.

Topology change management

When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Threat Defense device or the switch, adding an additional switch to form a VSS or vPC) you should disable the health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.

Operational considerations

Expect temporary, limited packet, or connection drop when adding a unit to an existing cluster or when reloading a unit.

- In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel, when the syslog server port is down and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the ASA cluster. These messages can result in some units of the ASA cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Default clustering settings

Review the default clustering configuration settings before deployment.

- The cLACP system ID is auto-generated, and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Firewall Mode

The firewall mode must match on all units.

High Availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Switches

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. In addition, we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR IPv4 MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.
- On the switch, we recommend that you use one of these EtherChannel load-balancing algorithms: **source-dest-ip** or **src-dst-mixed-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command).

  Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.
- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the cluster device to fixed:

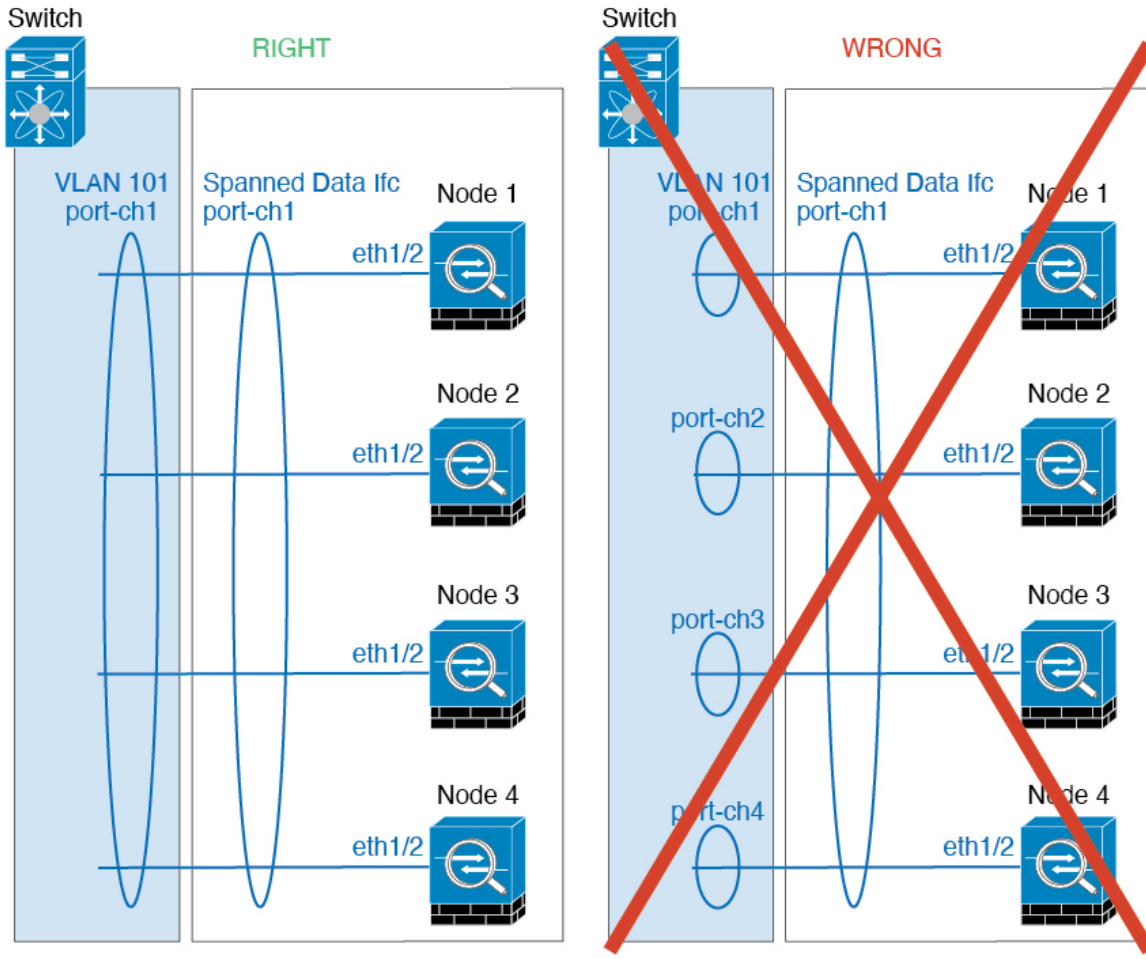
  router(config)# **port-channel** *id* **hash-distribution fixed**

  Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.
- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.
- Switches on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.
- Port-channel bundling downtime should not exceed the configured keepalive interval.
- You should disable the LACP Graceful Convergence feature on all cluster-facing EtherChannel interfaces for Cisco Nexus switches.

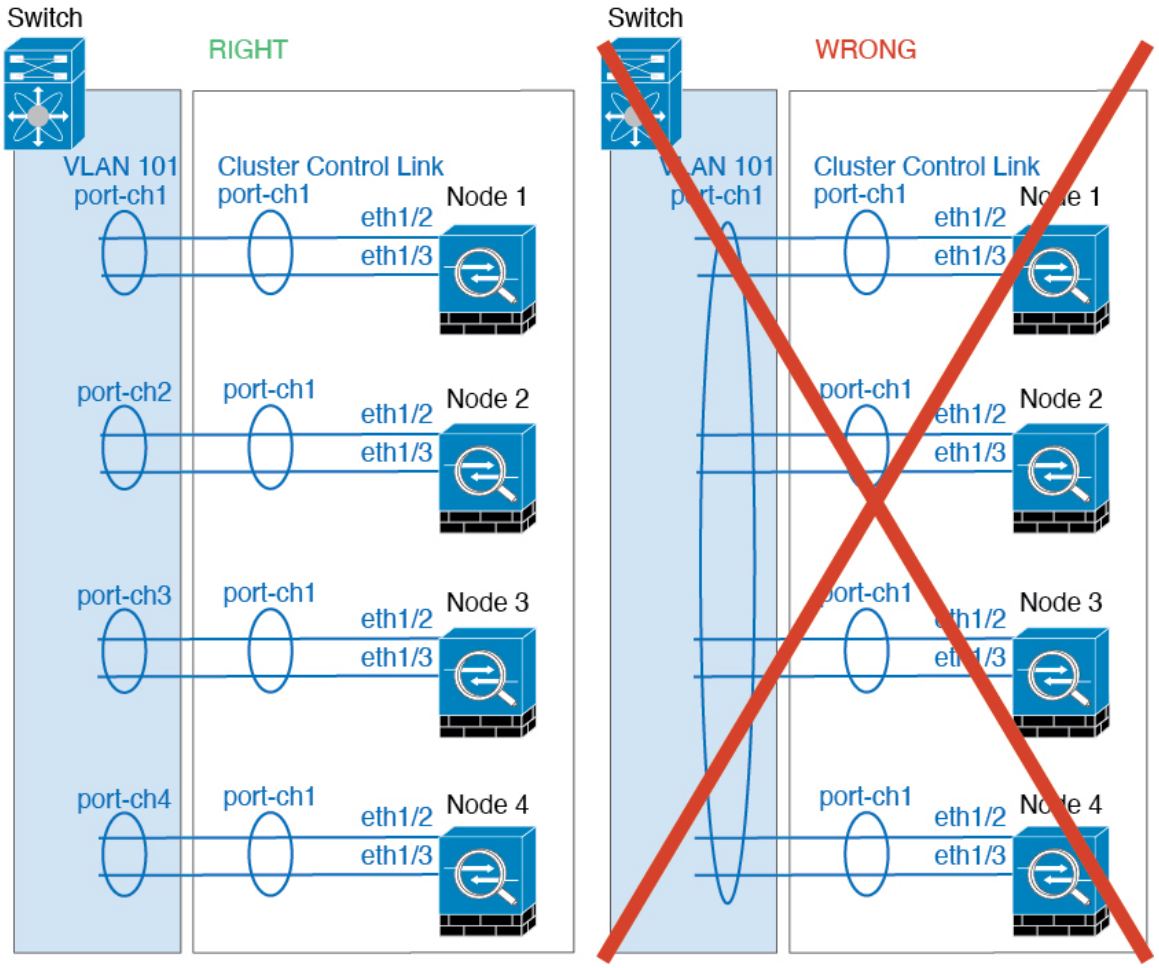
EtherChannels

- In Catalyst 3750-X Cisco IOS software versions earlier than 15.1(1)S2, the cluster unit did not support connecting an EtherChannel to a switch stack. With default switch settings, if the cluster unit EtherChannel is connected cross stack, and if the control unit switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.

- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.
  - Spanned EtherChannels—For cluster unit Spanned EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.



  - Device-local EtherChannels—For cluster unit Device-local EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels on the switch; do not combine multiple cluster unit EtherChannels into one EtherChannel on the switch.



Additional guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Threat Defense or the switch, adding an additional switch to form a VSS or vPC) you should disable the health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.
- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel, when the syslog server port is down and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the ASA cluster. These messages can result in some units of the ASA cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a

new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Defaults for clustering

- The cLACP system ID is auto-generated, and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Clustering configuration

To add a cluster to the Cloud-Delivered Firewall Management Center, add each node to the Cloud-Delivered Firewall Management Center as a standalone unit, configure interfaces on the unit you want to make the control node, and then form the cluster.

Cluster interfaces

You can configure data interfaces as either Spanned EtherChannels or as individual interfaces. All data interfaces in the cluster must be one type only. You cannot configure Ethernet 1/1 as a Spanned EtherChannel and configure Ethernet 1/2 as an Individual interface within the same cluster, for example.

For Spanned EtherChannels: You can use regular firewall interfaces or IPS-only interfaces (inline sets or passive interfaces). For Individual interfaces: IPS-only interfaces are not supported.

Each unit must also dedicate at least one hardware interface as the cluster control link.

Cluster control link

A cluster control link is a dedicated hardware interface that enables communication between cluster units and supports cluster operation coordination.

Each unit must dedicate at least one hardware interface as the cluster control link. We recommend using an EtherChannel for the cluster control link if available.

Cluster control link traffic

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

Cluster control link interfaces and network

You can use any physical interface or EtherChannel for the cluster control link. You cannot use a VLAN subinterface as the cluster control link. You also cannot use the Management interface.

Each cluster control link has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the cluster control link interfaces.



Note For a 2-member cluster, do not directly connect the cluster control link from one node to the other node. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit.

If you need to directly connect the units (for testing purposes, for example), configure and enable the cluster control link interface on both nodes before forming the cluster.

Size the cluster control link

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features.

For example:

- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.
- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.



Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster control link MTU ping testing

Use Cluster control link MTU ping testing to

- verify MTU compatibility between cluster nodes during the join process
- use ping-based fragmentation tests to confirm the network path supports proper packet handling, and
- allow nodes to join even if tests fail while providing diagnostic messages for troubleshooting MTU mismatches.

Control node ping test

When a node joins the cluster, the control node sends a ping with a payload size of twice the MTU. This process tests the network's ability to handle packet fragmentation because the underlying IP layer will fragment packets that exceed the MTU limit.

A successful ping confirms that the network path supports proper fragmentation and that the cluster control link can reliably process traffic at the configured MTU size.

If the ping fails, view these messages:

- **show cluster history**—`Event: CCL MTU test to unit name failed`
- Console WARNING—`WARNING: Unit name is not reachable in CCL jumbo frame ICMP test, please check cluster interface and switch MTU configuration`

Even if the ping fails, the node is allowed to join the cluster. In this case, you need to resolve the MTU mismatch as soon as possible.

Data node ping test

When a node joins the cluster, the joining node checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the initial ping fails, the node tries a ping using a smaller packet size (the MTU divided by 2, then by 4, then by 8) until a ping succeeds.

If the ping fails, view these messages:

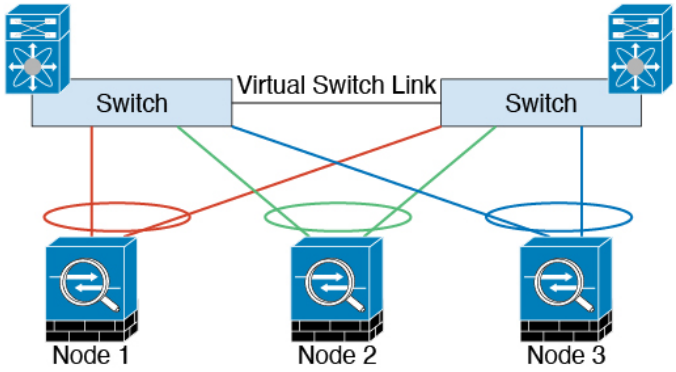
- **show cluster info trace**—`Warning: CCL MTU is configured to cfg_mtu_size. However CCL MTU test to unit name failed with size larger_test_size (passed with size smaller_test_size). Please check switch MTU configuration.`

  To easily view this WARNING, filter on the **show** output using **show cluster info trace | incl MTU**.
- **show cluster history**, Cloud-Delivered Firewall Management Center notification—`Warning: MTU mismatch detected on the CCL interface. Please ensure that the MTU setting on the connected switch matches the firewall's configured MTU (cfg_mtu_size).`
- Console WARNING—`WARNING: Unit name is not reachable in CCL jumbo frame ICMP test, please check cluster interface and switch MTU configuration`

Even if the ping fails, the node is allowed to join the cluster. In this case, you need to resolve the MTU mismatch as soon as possible.

Cluster control link redundancy

The diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS), Virtual Port Channel (vPC), StackWise, or StackWise Virtual environment. All links in the EtherChannel are active. When the switch is part of a redundant system, then you can connect firewall interfaces within the same EtherChannel to separate switches in the redundant system. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a Spanned EtherChannel.



Cluster control link reliability for inter-chassis clustering

To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

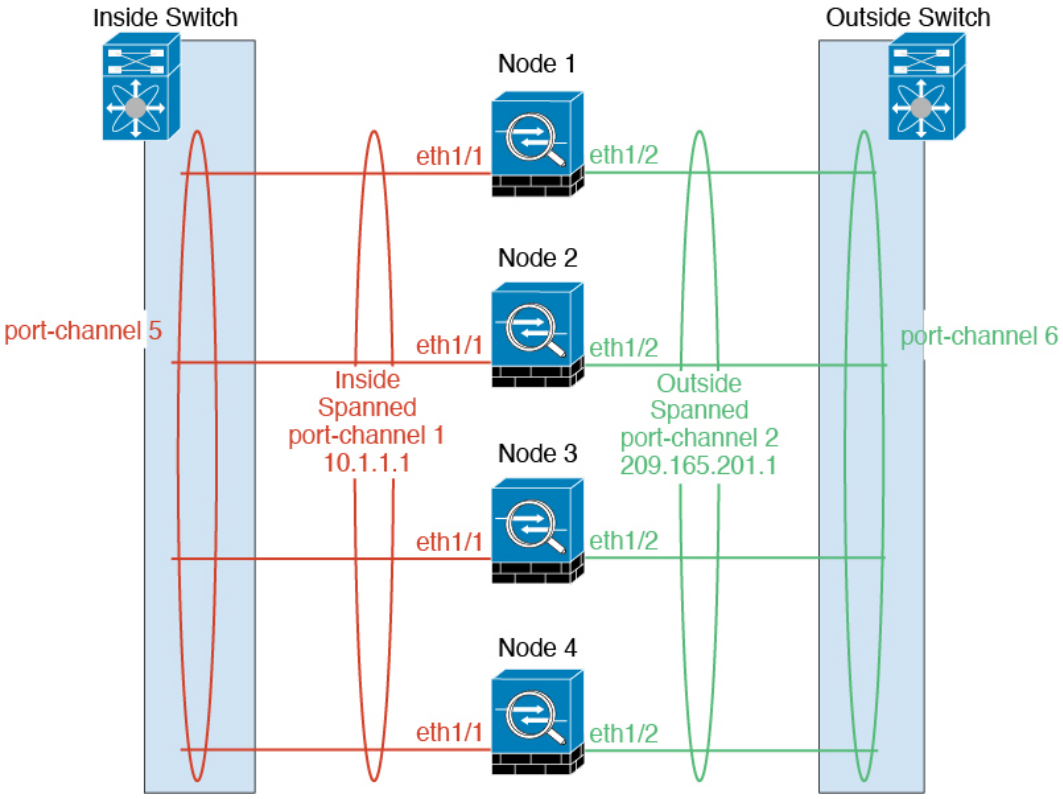
The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

Spanned EtherChannels (Recommended)

You can group one or more interfaces per chassis into an EtherChannel that spans all chassis in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel.

For regular firewall interfaces:A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address is assigned to the BVI, not to the bridge group member interface.

The EtherChannel inherently provides load balancing as part of basic operation.



Spanned EtherChannel benefits

The EtherChannel method of load-balancing is recommended over other methods for these benefits:

- Faster failure discovery.
- Faster convergence time. Individual interfaces rely on routing protocols to load-balance traffic, and routing protocols often have slow convergence during a link failure.
- Ease of configuration.

Guidelines for maximum throughput

To achieve maximum throughput, follow these guidelines:

- Use a symmetric load-balancing hash algorithm, meaning that packets from both directions will have the same hash and will be sent to the same Firewall Threat Defense in the Spanned EtherChannel. We recommend using the source and destination IP address (the default) or the source and destination port as the hashing algorithm.
- Use the same type of line cards when connecting the Firewall Threat Defenses to the switch so that hashing algorithms applied to all packets are the same.

Load balancing

The EtherChannel link is selected using a proprietary hash algorithm, based on source or destination IP addresses and TCP and UDP port numbers.

Note On the switch, we recommend that you use one of the following algorithms: **source-dest-IP** or **source-dest-IP-port** (see the Cisco Nexus OS or Cisco IOS **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the nodes in a cluster.

The number of links in the EtherChannel affects load balancing.

Symmetric load balancing is not always possible. If you configure NAT, then forward and return packets will have different IP addresses and/or ports. Return traffic will be sent to a different unit based on the hash, and the cluster will have to redirect most returning traffic to the correct unit.

EtherChannel redundancy

The EtherChannel has built-in redundancy. It monitors the line protocol status of all links. If one link fails, traffic is re-balanced between remaining links. If all links in the EtherChannel fail on a particular unit, but other units are still active, then the unit is removed from the cluster.

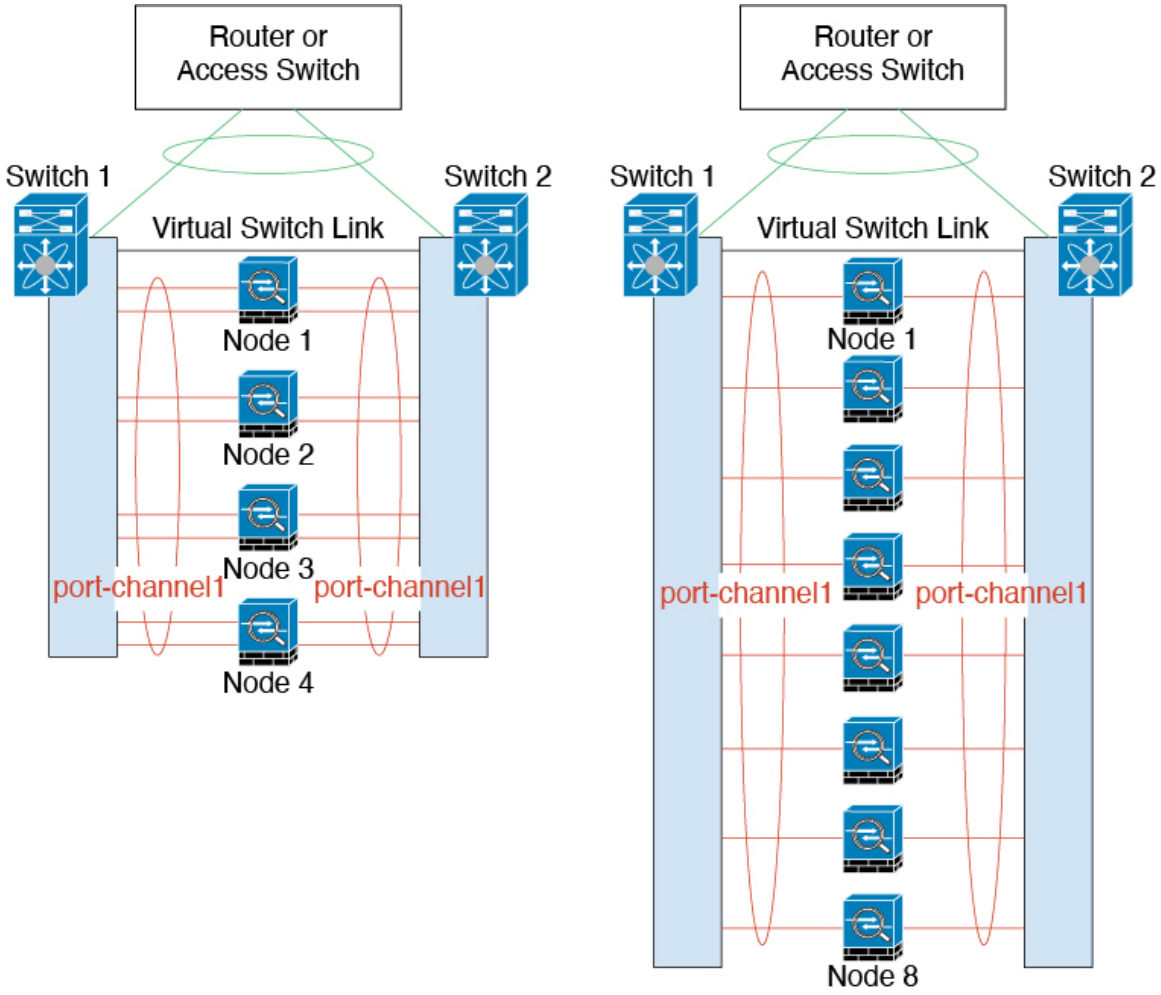
Connecting to a redundant switch system

You can include multiple interfaces per Firewall Threat Defense in the Spanned EtherChannel. Multiple interfaces per Firewall Threat Defense are especially useful for connecting to both switches in a VSS, vPC, StackWise, or StackWise Virtual system.

Depending on your switches, you can configure up to 32 active links in the spanned EtherChannel. This feature requires both switches in the vPC to support EtherChannels with 16 active links each (for example the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module).

For switches that support 8 active links in the EtherChannel, you can configure up to 16 active links in the spanned EtherChannel when connecting to two switches in a redundant system.

This figure shows a 16-active-link spanned EtherChannel in a 4-node cluster and an 8-node cluster.



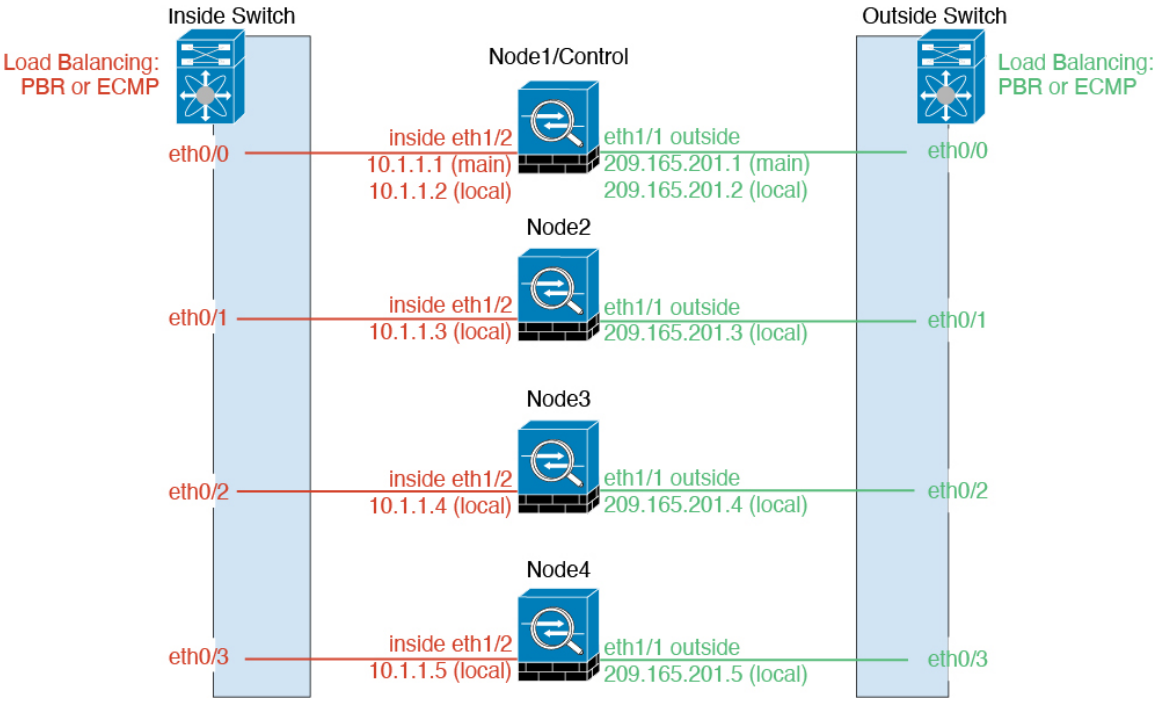
Individual interfaces

Individual interfaces are normal routed interfaces, each with their own *local IP address* used for routing. The *main cluster IP address* for each interface is a fixed address that always belongs to the control node. When the control node changes, the Main cluster IP address moves to the new control node, so management of the cluster continues seamlessly.

IPS-only interfaces (inline sets and passive interfaces) are not supported as Individual interfaces.

Because interface configuration must be configured only on the control node, you configure a pool of IP addresses to be used for a given interface on the cluster nodes, including one for the control node.

Load balancing must be configured separately on the upstream switch.



Policy-based routing

When using individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Policy-Based Routing (PBR).

We recommend this method if you are already using PBR, and want to take advantage of your existing infrastructure.

PBR makes routing decisions based on a route map and ACL. You must manually divide traffic between all Firewall Threat Defenses in a cluster. Because PBR is static, it may not achieve the optimum load balancing result at all times. To achieve the best performance, we recommend that you configure the PBR policy so that forward and return packets of a connection are directed to the same Firewall Threat Defense.

For example, if you have a Cisco router, redundancy can be achieved by using Cisco IOS PBR with Object Tracking. Cisco IOS Object Tracking monitors each Firewall Threat Defense using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular Firewall Threat Defense.

Equal-cost multi-path routing

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Equal-Cost Multi-Path (ECMP) routing.

We recommend this method if you are already using ECMP, and want to take advantage of your existing infrastructure.

ECMP routing considerations

ECMP routing can forward packets over multiple best paths that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses and/or source and destination ports can be used to send a packet to one of the next hops.

If you use static routes for ECMP routing, then the Firewall Threat Defense failure can cause problems; the route continues to be used, and traffic to the failed Firewall Threat Defense will be lost.

If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each Firewall Threat Defense to participate in dynamic routing.

Cisco Intelligent Traffic Director (Routed Firewall Mode Only)

Cisco Intelligent Traffic Director (ITD) is a high-speed hardware load-balancing solution that

- supports Nexus 5000, 6000, 7000, and 9000 switch series
- fully covers the functional capabilities of traditional PBR with a simplified configuration workflow, and
- offers multiple additional features for more granular load distribution.

ITD capabilities and configuration

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address.

ITD provides these load-balancing capabilities:

- IP stickiness
- Consistent hashing for bi-directional flow symmetry
- Virtual IP addressing
- Health monitoring
- Sophisticated failure handling policies with N+M redundancy
- Weighted load-balancing
- Application IP SLA probes including DNS

Due to the dynamic nature of load-balancing, ITD achieves a more even traffic distribution across all cluster nodes as compared to PBR. To achieve bi-directional flow symmetry, configure ITD such that forward and return packets of a connection are directed to the same Firewall Threat Defense.

For more details, refer to Cisco Intelligent Traffic Director Deployment Guide with Cisco ASA.

Cable and add devices to the Cloud-Delivered Firewall Management Center

Before configuring clustering, you need to prepare your devices. In particular, the cluster will not come up unless all nodes can communicate over the cluster control link. Therefore, before you form the cluster, the cluster control link must be ready to go.

Procedure

Step 1 Cable the cluster control link network, management network, and data networks.

Step 2 Configure the upstream and downstream equipment.

a) For the cluster control link network, set the MTU to be at least 100 bytes higher than the data interface MTU.

   By default, the data interface MTU is 1500 bytes, so the cluster control link MTU on the cluster node will be set to 1600 bytes. If you use higher MTUs on your data interfaces, increase the cluster control link MTU on connecting switches accordingly.

b) Configure cluster control link interfaces on upstream and downstream equipment, including for an optional EtherChannel.

   For more information, refer to Cluster control link interfaces and network, on page 732 for cluster control link requirements.

c) Configure data interfaces on upstream and downstream equipment, including Spanned EtherChannels.

   For more information, refer to Cluster interfaces, on page 731 for information about how to cable Spanned EtherChannels.

Step 3 Add each node to the Cloud-Delivered Firewall Management Center as a standalone device in the same domain and group.

You can create a cluster with a single device, and then add more nodes later. The initial settings (licensing, access control policy) that you set when you add a device will be inherited by all cluster nodes from the control node. You will choose the control node when forming the cluster.

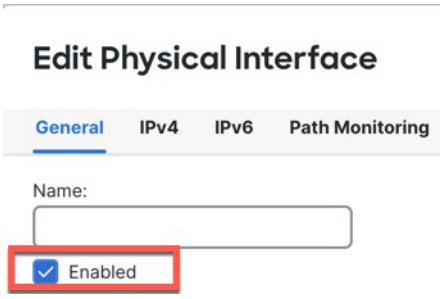
Step 4 Enable the cluster control link on the device you want to be the control node.

When you add the other nodes, they will inherit the cluster control link configuration.

Note
Do *not* configure the name or IP addressing for the cluster control link. The MTU of the cluster control link interface is automatically set to 100 bytes more than the highest data interface MTU when you form the cluster, so you do not need to set it now. However, we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation. If the MTU is set in this range when you add the cluster, we recommend returning to the **Interfaces** page and manually increasing it above 8362. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the initial ping fails, the node tries a ping using a smaller packet size (the MTU divided by 2, then by 4, then by 8) until a ping succeeds. A notification is generated so you can fix the MTU mismatch on connecting switches and try again.

a) On the device you want to be the control node, choose **Devices** > **Device Management**, and click **Edit** ().
b) Click **Interfaces**.
c) Enable the interface. If you are going to use an EtherChannel for the cluster control link, enable all members.For more information, refer to Enable the physical interface and configure Ethernet settings, on page 1003.

Figure 104: Enable the cluster control link interface



d) (Optional) Add an EtherChannel.For more information, refer to Configure an EtherChannel, on page 1011.

We recommend using the On mode for cluster control link member interfaces to reduce unnecessary traffic on the cluster control link (Active mode is the default). The cluster control link does not need the overhead of LACP traffic because it is an isolated, stable network. **Note:** We recommend setting *data* EtherChannels to Active mode.

e) Click **Save** and then **Deploy** to deploy the interface changes to the control node.

Create a Cluster

Form a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Device Management**, and then choose **Add** > **Cluster**.

The **Add Cluster Wizard** appears.

Figure 105: Add Cluster Wizard

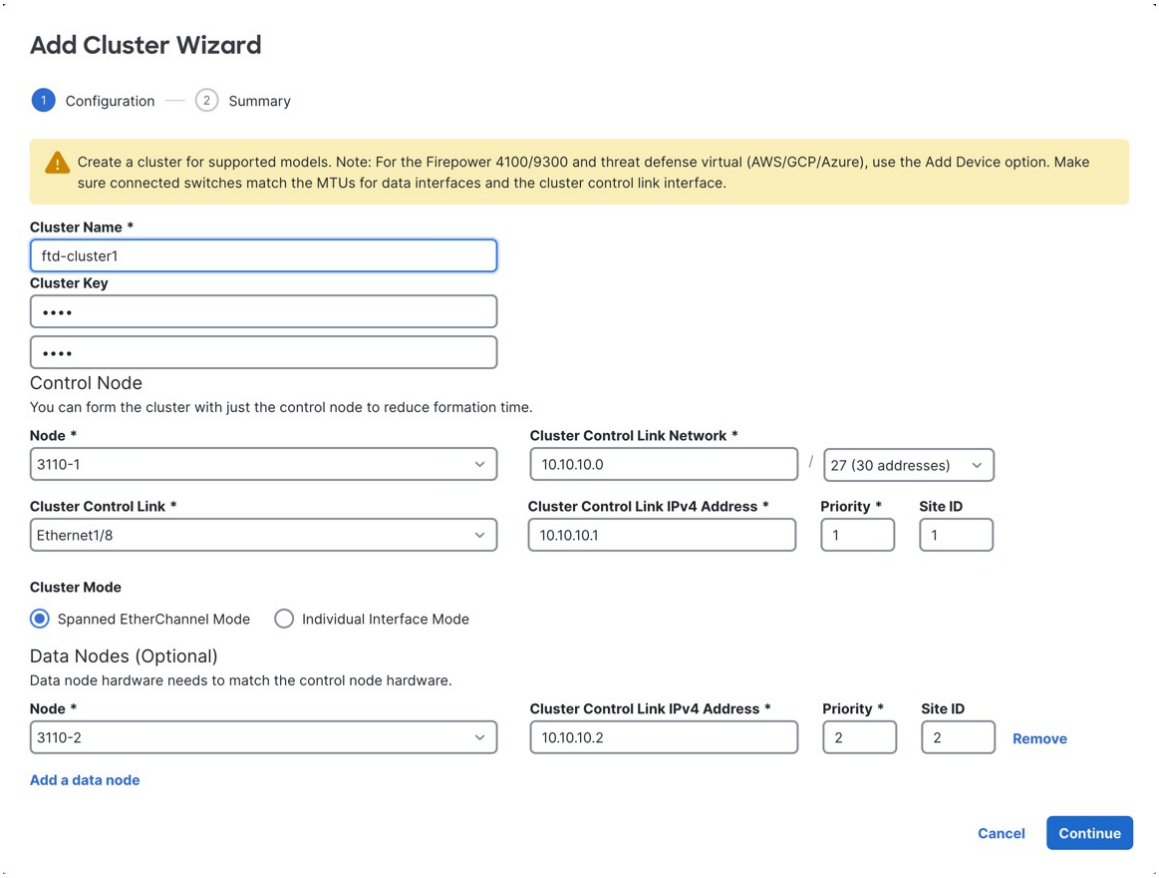


Figure 106:

Step 2 Specify a **Cluster Name** and an authentication **Cluster Key** for control traffic.

- **Cluster Name**—An ASCII string of 1 to 38 characters. The cluster name can contain only letters, digits, underscores (_), and hyphens (-).
- **Cluster Key**—An ASCII string of 1 to 63 characters. The cluster key cannot be a single digit and must not contain spaces. The **Cluster Key** value is used to generate the encryption key. This encryption does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in clear text.

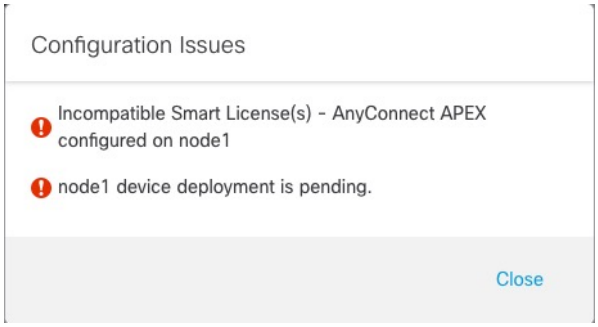
Step 3 For the **Control Node**, configure these parameters:

- **Node**—Choose the device that you want to be the control node initially. When the Cloud-Delivered Firewall Management Center forms the cluster, it will add this node to the cluster first so it will be the control node.

  Note

  If you see an **Error** ( ) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation. For example:

Figure 107: Configuration Issues



To resolve the above issues, remove the unsupported VPN license and deploy pending configuration changes to the device.

- **Cluster Control Link Network**—Specify an IPv4 subnet; IPv6 is not supported for this interface. Specify a **24**, **25**, **26**, or **27** subnet.
- **Cluster Control Link**—Choose the physical interface or EtherChannel you want to use for the cluster control link.

  Note
  The MTU of the cluster control link interface is automatically set to 100 bytes more than the highest data interface MTU; by default, the MTU is 1600 bytes. We do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation. If the MTU is set in this range when you add the cluster, we recommend increasing the MTU above 8362. Choose **Devices** > **Device Management** and then click **Interfaces**.

  Make sure you configure switches connected to the cluster control link to the correct (higher) MTU; otherwise, cluster formation can fail.

- **Cluster Control Link IPv4 Address**—This field will be auto-populated with the first address on the cluster control link network. You can edit the host address if desired.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority. Even if you set the priority to be lower than other nodes, this node will still be the control node when the cluster is first formed.
- **Site ID**—(FlexConfig feature) Enter the site ID for this node between 1 and 8. A value of 0 disables inter-site clustering. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FlexConfig feature.

Step 4 For the **Cluster Mode**, choose **Spanned EtherChannel Mode** or **Individual Interface Mode** .

Step 5 For **Data Nodes (Optional)**, click **Add a data node** to add a node to the cluster.

You can form the cluster with only the control node for faster cluster formation, or you can add all nodes now. Set the following for each data node:

- **Node**—Choose the device that you want to add.

  Note

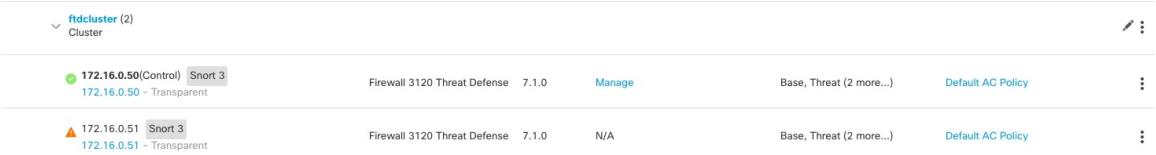
If you see an **Error** ( ) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation.

- **Cluster Control Link IPv4 Address**—This field will be auto-populated with the next address on the cluster control link network. You can edit the host address if desired.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority.
- **Site ID**—(FlexConfig feature) Enter the site ID for this node between 1 and 8. A value of 0 disables inter-site clustering. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FlexConfig feature.

Step 6 Click **Continue**. Review the **Summary**, and then click **Save**.

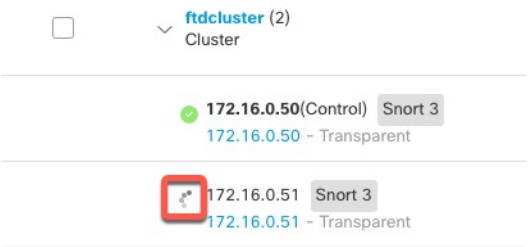
The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster nodes.

Figure 108: Cluster Management

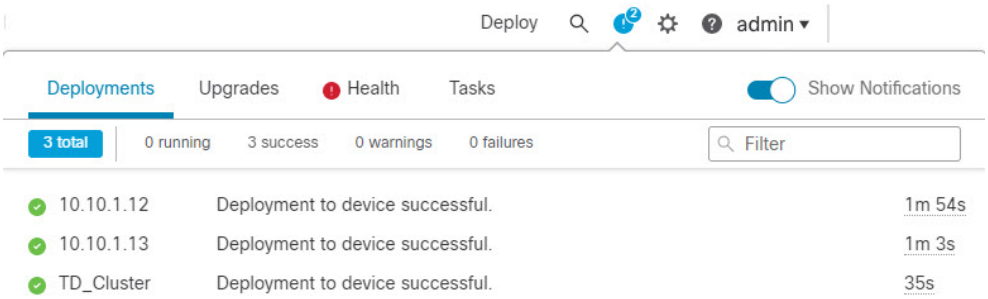


A node that is currently registering shows the loading icon.

Figure 109: Node Registration



You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**. The Cloud-Delivered Firewall Management Center updates the Cluster Registration task as each node registers.

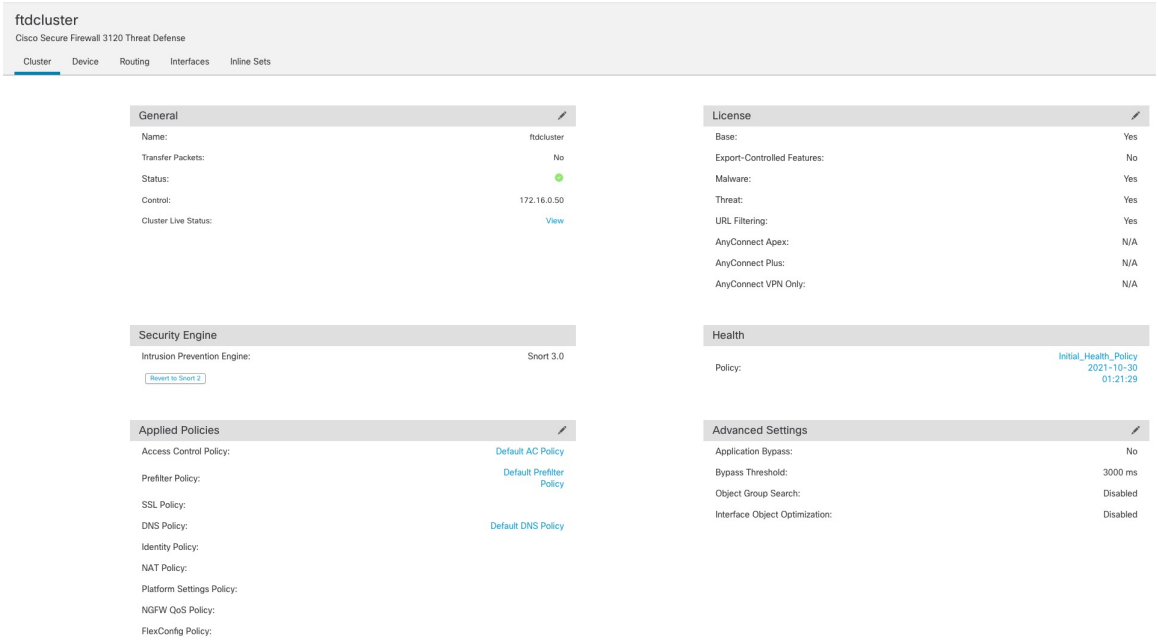


Step 7 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

Most configuration can be applied to the cluster as a whole, and not nodes in the cluster. For example, you can change the display name per node, but you can only configure interfaces for the whole cluster.

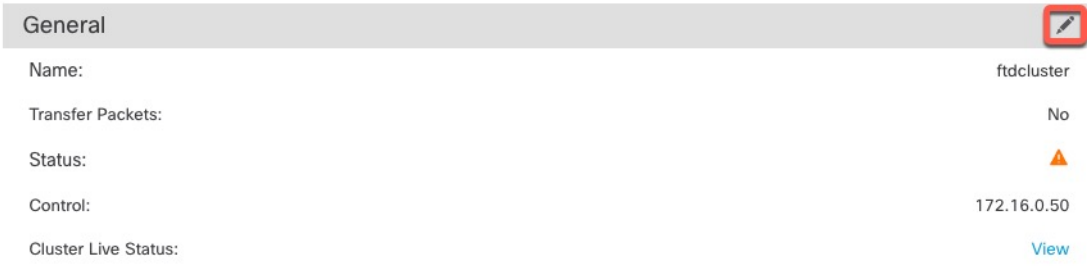
Step 8 On the **Devices** > **Device Management** and then choose **Add**, **Cluster** screen, you see **General** and other settings for the cluster.

Figure 110: Cluster Settings

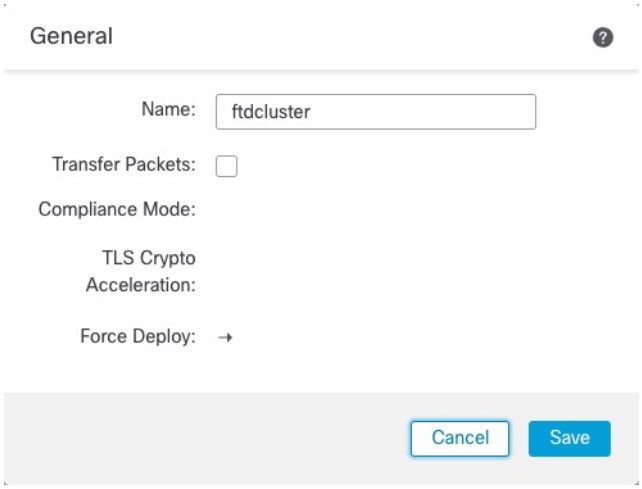


See these cluster-specific items in the **General** area:

- **General > Name**—Change the cluster display name by clicking the **Edit** (✎).



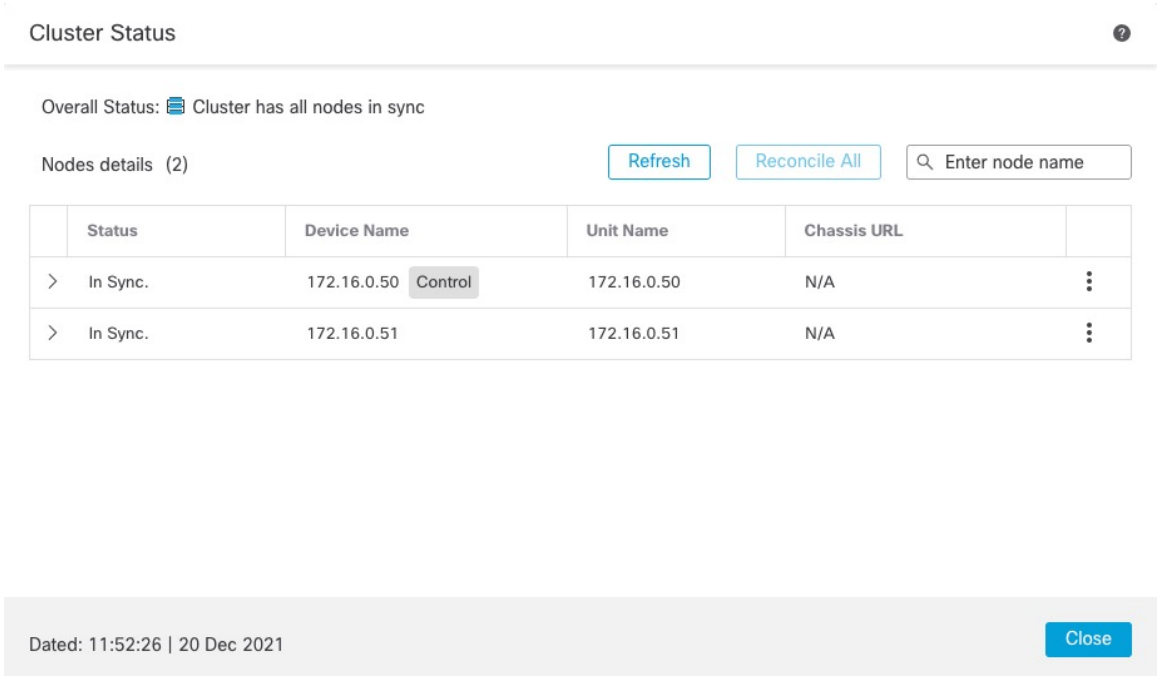
Then set the **Name** field.



- **General > Cluster Live Status**—Click the **View** link to open the **Cluster Status** dialog box.



The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile All**.



- **General > Troubleshoot**—You can generate and download troubleshooting logs, and you can view cluster CLIs. Refer Troubleshoot the cluster, on page 777 .

Figure 111: Troubleshoot



Step 9 On the **Devices** > **Device Management** and then choose **Add**, **Device**, you can choose each member in the cluster from the top right drop-down menu and configure these settings.

Figure 112: Device Settings

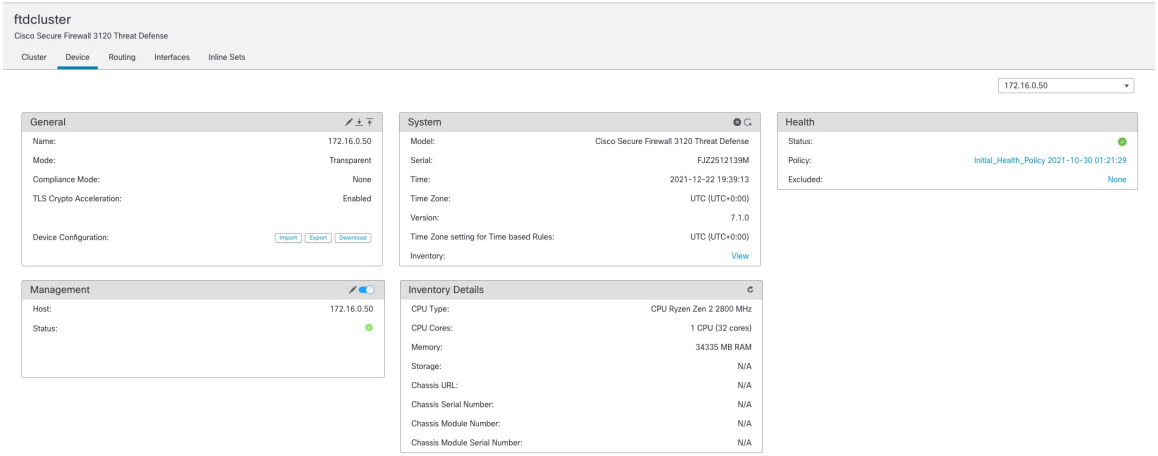
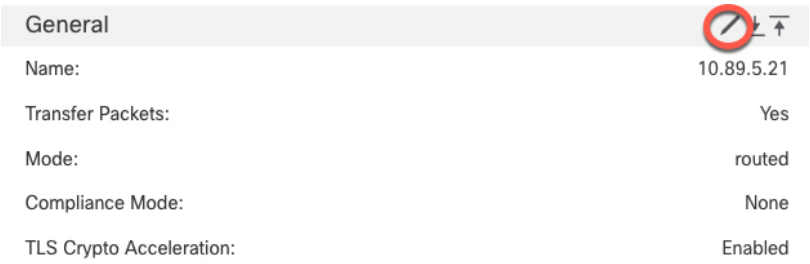


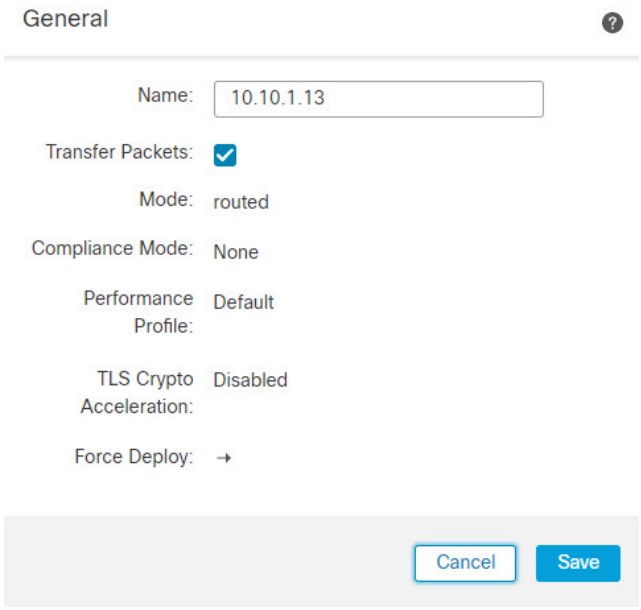
Figure 113: Choose Node



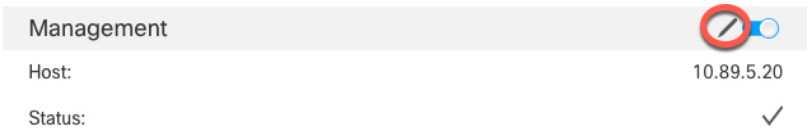
- **General > Name**—Change the cluster member display name by clicking the **Edit** ( ) .



Then set the **Name** field.



- **Management > Host** —If you change the management IP address in the device configuration, you must match the new address in the Cloud-Delivered Firewall Management Center so that it can reach the device on the network. First disable the connection, edit the **Host** address in the **Management** area, then re-enable the connection.



Configure interfaces

You can configure data interfaces as either Spanned EtherChannels or as Individual interfaces. Each method uses a different load-balancing mechanism. You cannot configure both types in the same configuration.

Configure spanned EtherChannels

Configure data interfaces as Spanned EtherChannels.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) next to the cluster.

Step 2 Click **Interfaces**.

Step 3 Configure Spanned EtherChannel data interfaces.

a) Configure one or more EtherChannels. See Configure an EtherChannel, on page 1011.

You can include one or more member interfaces in the EtherChannel. Because this EtherChannel is spanned across all of the nodes, you only need one member interface per node; however, for greater throughput and redundancy, multiple members are recommended.

b) (Optional) For regular firewall interfaces, configure VLAN subinterfaces on the EtherChannel. The rest of this procedure applies to the subinterfaces. See Add a subinterface, on page 1060.

c) Click **Edit** ( ) for the EtherChannel interface.

d) Configure the name and other parameters. For regular firewall interfaces, see Configure routed mode interfaces, on page 1080 or, for transparent mode, Bridge group interfaces, on page 1086. For IPS-only interfaces, see Inline Sets and Passive Interfaces, on page 1119.

- If the cluster control link interface MTU is not at least 100 bytes higher than the data interface MTU, you will see an error that you must reduce the MTU of the data interface. By default, the cluster control link MTU is 1600 bytes. If you want to increase the MTU of data interfaces, first increase the cluster control link MTU. Note that we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.
- For routed mode, DHCP, PPPoE, IPv6 autoconfig and manual link-local addresses are not supported. For point-to-point connections, you can specify a 31-bit subnet mask (255.255.255.254). In this case, no IP addresses are reserved for the network or broadcast addresses.

e) Set a unique, manual global MAC address for the EtherChannel. Click **Advanced**, and in the **Active MAC Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

Do not set the **Standby MAC Address**; it is ignored.

You must configure a unique MAC address not currently in use on your network for a Spanned EtherChannel to avoid potential network connectivity problems. With a manually-configured MAC address, the MAC address stays with the current control unit. If you do not configure a MAC address, then if the control unit changes, the new control unit uses a new MAC address for the interface, which can cause a temporary network outage.

f) Click **OK**. Repeat the above steps for other data interfaces.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure individual interfaces

Individual interfaces are normal routed interfaces, each with their own IP address taken from a pool of IP addresses. The Main cluster IP address is a fixed address for the cluster that always belongs to the control node.

Individual management interfaces let you SSH directly to each unit if necessary, while a Spanned EtherChannel interface only allows connection to the control node.

IPS-only interfaces (inline sets or passive interfaces) are not supported for Individual interfaces.

Before you begin

- You must be in Individual interface mode.
- Individual interfaces require you to configure load balancing on neighbor devices. External load balancing is not required for the management interface.
- (Optional) Configure the interface as a device-local EtherChannel interface, and/or configure subinterfaces.

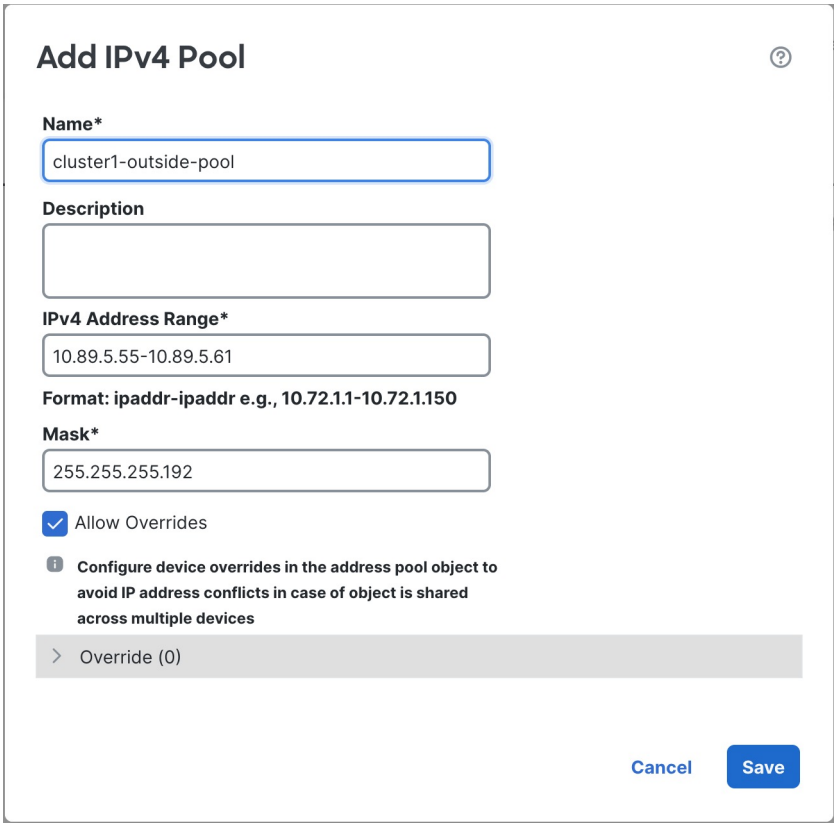
  For an EtherChannel, this EtherChannel is local to the unit, and is not a Spanned EtherChannel.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Address Pools** to add an IPv4 and/or IPv6 address pool.For more information, refer to Configure address pools, on page 2954.

Include at least as many addresses as there are units in the cluster. The main IP address is not a part of this pool, but needs to be on the same network. You cannot determine the exact Local address assigned to each unit in advance.

Figure 114: Add address pool



Note

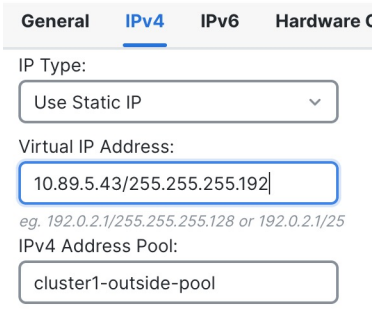
Although not common, if you want to set the MAC addresses manually, you can also add a MAC address pool object.

Step 2 On **Devices** > **Device Management** and then choose **Interfaces**, click **Edit** (✎) for the interface you want to configure.

Step 3 On the **IPv4** page, enter the **Virtual IP Address** and mask. This main ("virtual") IP address is a fixed address for the cluster, and always belongs to the control node.

DHCP and PPPoE are not supported. For point-to-point connections, you can specify a 31-bit subnet mask (255.255.255.254). In this case, no IP addresses are reserved for the network or broadcast addresses.

Figure 115: IPv4 page



Step 4 From the **IPv4 Address Pool** drop-down list, choose the address pool you created.

Step 5 On **IPv6 > Basic**, from the **IPv6 Address Pool** drop-down list, choose the address pool you created.

IPv6 autoconfig and manual link-local addresses are not supported.

Step 6 Configure other interface settings as normal.

To set the MAC addresses manually, you can select the MAC address pool from the interface's **Advanced** page.

Note
If the cluster control link interface MTU is not at least 100 bytes higher than the data interface MTU, you will see an error that you must reduce the MTU of the data interface. By default, the cluster control link MTU is 1600 bytes. If you want to increase the MTU of data interfaces, first increase the cluster control link MTU. Note that we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.

Configure cluster health monitor settings

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings described in the table below.

Figure 116: Cluster Health Monitor Settings

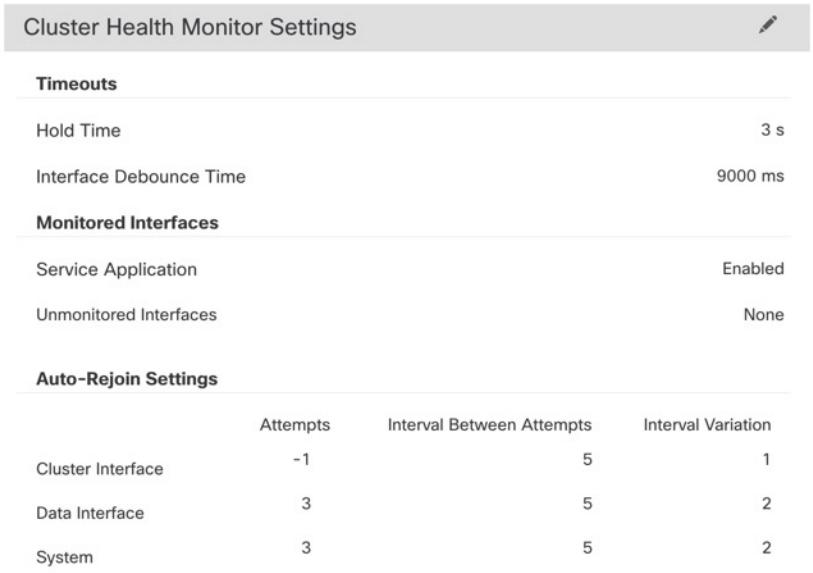


Table 50: Cluster Health Monitor Settings Section Table Fields

Field	Description
Timeouts	
Hold Time	Between .3 and 45 seconds; The default is 3 seconds. To determine node system health, the cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the hold time period, the peer node is considered unresponsive or dead.
Interface Debounce Time	Between 300 and 9000 ms. The default is 500 ms. The interface debounce time is the amount of time before the node considers an interface to be failed, and the node is removed from the cluster.
Monitored Interfaces	The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster.
Service Application	Shows whether the Snort and disk-full processes are monitored.
Unmonitored Interfaces	Shows unmonitored interfaces.
Auto-Rejoin Settings	
Cluster Interface	Shows the auto-rejoin settings after a cluster control link failure.

Field	Description
Attempts	Between -1 and 65535. The default is -1 (unlimited). Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 1x the interval duration. Defines if the interval duration increases at each attempt.
Data Interfaces	Shows the auto-rejoin settings after a data interface failure.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.
System	Shows the auto-rejoin settings after internal errors. Internal failures include: application sync timeout; inconsistent application statuses; and so on.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.

Note If you disable the system health check, fields that do not apply when the system health check is disabled will not show.

You can change these settings from this section.

You can monitor any port-channel ID, single physical interface ID, as well as the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Devices** > **Device Management**.

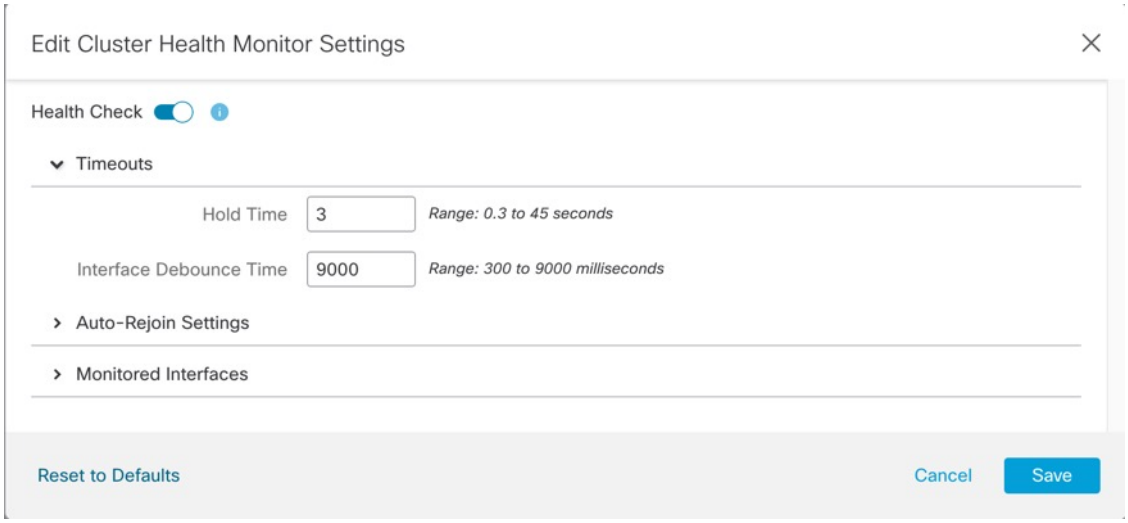
Step 2 Next to the cluster you want to modify, click **Edit** ( ).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** ( ).

Step 5 Disable the system health check by clicking the **Health Check** slider .

Figure 117: Disable the System Health Check



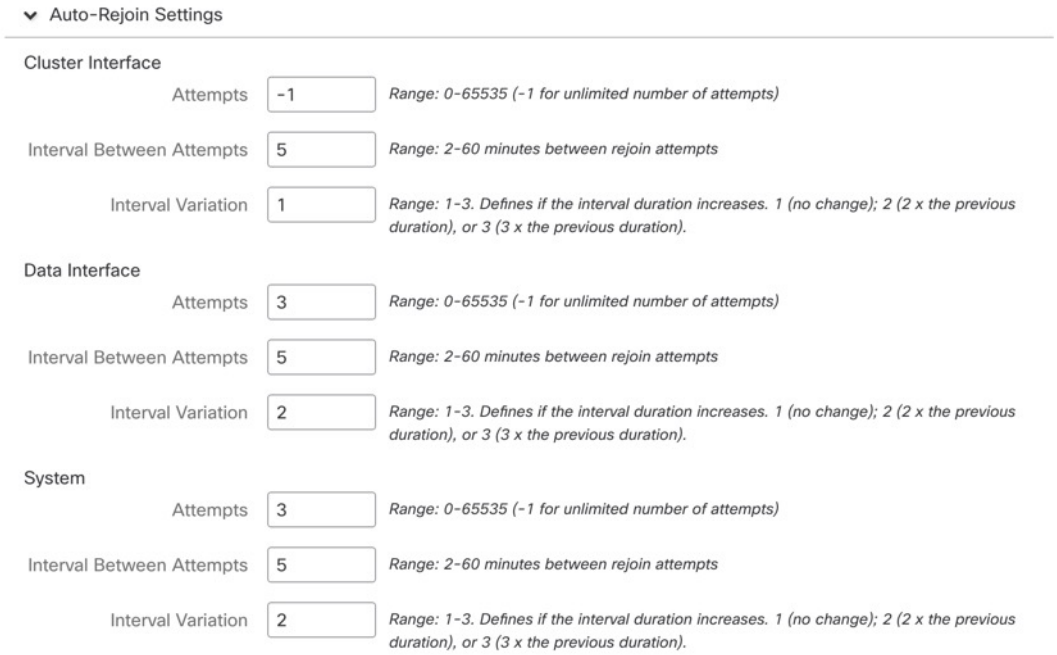
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the number of milliseconds specified before marking the interface as failed, and the node is removed from the cluster. In the case of an EtherChannel that transitions from a down state to an up state (for example, the switch reloaded, or the switch enabled an EtherChannel), a longer debounce time can prevent the interface from appearing to be failed on a cluster node just because another cluster node was faster at bundling the ports.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 118: Configure Auto-Rejoin Settings

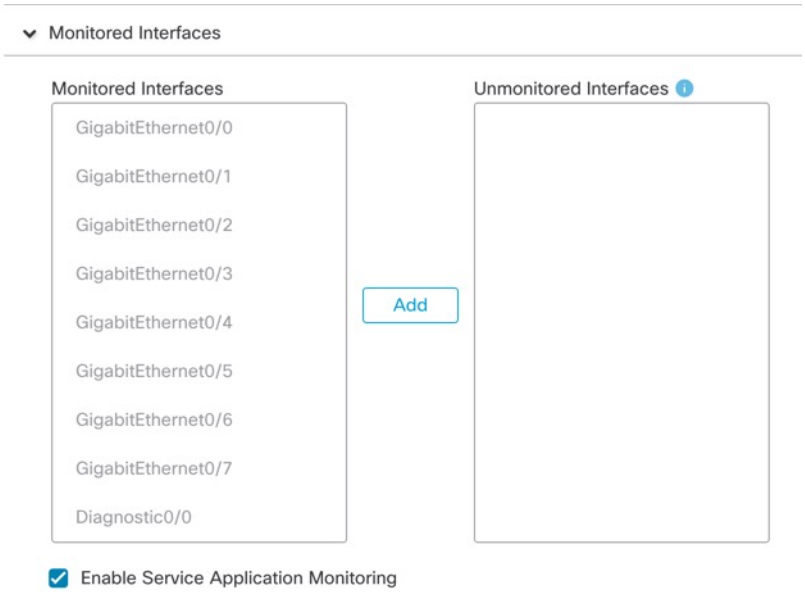


Set the following values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 119: Configure Monitored Interfaces



The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster. Health check is enabled by default for all interfaces and for the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save** and deploy the configuration changes.

Configure distributed Site-to-Site VPN

By default, the cluster uses centralized site-to-site VPN mode. To take advantage of the scalability of clustering, you can enable distributed site-to-site VPN mode.

Distributed site-to-site VPN

In distributed mode, site-to-site IPsec IKEv2 VPN connections are distributed across nodes of a cluster. Distributing VPN connections across the nodes of a cluster allows both the capacity and throughput of the cluster to be fully utilized, significantly scaling VPN support beyond centralized VPN capabilities.

Distributed VPN connection roles

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation. When running in distributed VPN mode, the following roles are assigned to the cluster nodes:

- **Active Session Owner**—The node that initially receives the connection or that has transitioned a backup session to an active session. The owner maintains the state and processes packets for the complete session, including the IKE and IPsec tunnels and all traffic associated with them.
- **Backup Session Owner**—The node that is handling the backup session for an existing active session. If the active session owner fails, the backup session owner becomes the active session owner, and a new backup session is established on a different node.
- **Forwarder**—If traffic associated with a VPN session is sent to a node that does not own the VPN session, that node will use the cluster control link to forward the traffic to the node that owns the VPN session.
- **Orchestrator**—The orchestrator (always the control node of the cluster) is responsible for calculating which sessions will move and where they will move to when executing an Active Session Redistribution (ASR). It sends a request to the owner node X to move N sessions to node Y. Node X will respond back to the orchestrator when complete, specifying how many sessions it was able to move.

Distributed VPN session characteristics

Distributed site-to-site VPN sessions have the following characteristics. Otherwise, VPN connections behave as they normally do if not in a cluster.

- VPN sessions are distributed across the cluster at the session level. Meaning the same cluster node handles the IKE and IPsec tunnels and all their traffic for a VPN connection. If VPN session traffic is sent to a cluster node that does not own that VPN session, traffic is forwarded to the cluster node that owns the VPN session.
- VPN sessions have a Session ID that is unique across the cluster. Using the session ID, traffic is validated, forwarding decisions are made, and IKE negotiation is completed.
- In a site-to-site VPN hub and spoke configuration, when clients connect through the cluster (called hair-pinning), the session traffic flowing in and the session traffic flowing out may be on different cluster nodes.

Distributed VPN handling of cluster events

This table lists how distributed VPN responds to different cluster events including node failures, node inactivation, and node joins.

Event	Distributed VPN
Node failure	For all active sessions on this failed node, the backup sessions (on another node) become active, and backup sessions are reallocated on another node.
Inactivate a cluster node	For all active sessions on the cluster node being inactivated, backup sessions (on another node) become active and reallocate backup sessions on another node according to the backup strategy.

Event	Distributed VPN
Cluster node join	If the VPN cluster mode on the new node is not set to distributed, the control node will request a mode change. After the VPN mode is compatible, the cluster node will be assigned active and backup sessions in the flow of normal operations.

IPsec IKEv2 modifications

IKEv2 is modified while in distributed site-to-site VPN mode in the following ways:

- An identity is used in place of IP/port tuples. This allows for proper forwarding decisions on the packets, and cleanup of previous connections that may be on other cluster members.
- The (SPI) identifiers that identify a single IKEv2 session are locally generated, random 8-byte values that are unique across the cluster. An SPI embeds a time stamp and a cluster node ID. Upon receipt of an IKE negotiation packet, if the time stamp or cluster node ID check fails, the packet is dropped and a message is logged indicating the reason.
- IKEv2 processing has been modified to prevent NAT-T negotiations from failing by being split across cluster members. A new ASP classify domain, **cluster_isakmp_redirect**, and rules are added when IKEv2 is enabled on an interface.

Certificate Management Protocol version 2

The Certificate Management Protocol version 2 (CMPv2) ID certificate and key pairs are synchronized across the cluster nodes. However, only the control node in the cluster automatically renews and rekeys the CMPv2 certificate. The control node synchronizes these new ID certificates and keys to all cluster nodes on a renewal. In this way, all nodes in the cluster utilize the CMPv2 certificates for authentication, and also any node is capable of taking over as the control node.

License for distributed site-to-site VPN

A Carrier license on each member of the cluster is required for distributed site-to-site VPN.

Prerequisites for distributed site-to-site VPN

Review these prerequisites before implementing distributed site-to-site VPN.

Model support

- Secure Firewall 4200

Cluster requirements

- Spanned EtherChannel mode.
- Routed firewall mode.

Maximum VPN sessions

Each VPN connection requires two sessions, one for the active session and one for the backup session. The maximum VPN session capacity of the cluster can be no more than half of the listed capacity due to using two licenses for each session.

Table 51: Maximum VPN sessions

Model	Maximum VPN Sessions
4215	10,000
4225	12,500
4245	15,000

Guidelines for distributed site-to-site VPN

Firewall mode

Distributed site-to-site VPN is supported in routed mode only.

Unsupported inspections

These types of inspections are not supported or are disabled in distributed site-to-site VPN mode:

- CTIQBE and DCERPC
- H323, H225, and RAS
- IPsec pass-through
- MGCP, MMP and NetBIOS
- PPTP and RADIUS
- RSH and RTSP
- SCCP (Skinny) and SUNRPC
- TFTP
- WAAS, WCCP, XDMCP

Additional Guidelines

- Only IKEv2 IPsec site-to-site VPN is supported in distributed site-to-site VPN mode. IKEv1 is not supported. IKEv1 site-to-site is supported in centralized VPN mode.
- Inter-site clustering is not supported.
- To view FlexConfig features that are also not supported with clustering, for example many inspections, refer to ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Cloud-Delivered Firewall Management Center GUI. For more information, refer to FlexConfig policies, on page 1309.

Enable distributed site-to-site VPN

Enable distributed site-to-site VPN to take advantage of the scalability of clustering for VPN sessions. This procedure requires CLI access to the control node.

Before you begin

- Configure site-to-site VPN according to Site-to-Site VPN, on page 1725.
- Apply the Carrier license to the control node.

Procedure

Step 1 Add a Spanned-EtherChannel-mode cluster according to Create a Cluster, on page 741, including all data nodes.

Step 2 After the cluster is formed and stable, remove each of the data nodes according to Disable clustering, on page 766, and then remove the control node.

Clustering needs to be disabled to change the VPN mode. The bootstrap configuration remains intact, as well as the last configuration synched from the control node, so that you can later re-add the nodes without losing your configuration.

Step 3 Connect to the control node CLI. Refer to the 'Log Into the Command-Line Interface on the Device' section in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

View the cluster on **Devices** > **Device Management** to see which device is the control node.

Step 4 Enable distributed site-to-site VPN on the control node.

cluster VPN-mode distributed

To disable distributed site-to-site VPN, use the **cluster VPN-mode centralized** command.

Example:

```
> cluster vpn-mode distributed
Cryptochecksum: ce4b0bbd 6b9252a5 7e19463d e179067d

5778 bytes copied in 0.70 secs
>
```

Step 5 In the Cloud-Delivered Firewall Management Center, reenable clustering on the control node and then for each data node. For more information, refer to Rejoin the cluster, on page 767.

The VPN mode is synched to the data nodes.

Redistribute distributed site-to-site VPN sessions

Active session redistribution redistributes the active VPN session load across the cluster nodes.

Due to the dynamic nature of beginning and ending sessions, active session redistribution is a best effort balancing of the sessions across all cluster nodes. Repeated redistribution actions will optimize the balance.

Redistribution can be run at any time, should be run after any topology change in the cluster, and is recommended after a new node joins the cluster.

To move a session, the backup session becomes the active one and another node is selected to host a new backup session. Moving sessions is dependent on the location of the active session's backup and the number of active sessions already on that particular backup node. If the backup session node is unable to host the active session for some reason, the original node remains owner of the session.

Before you begin

- Enable system logs if you would like to monitor redistribution activity.
- Enable CLI access to the control node.

Procedure

Step 1 Connect to the control node CLI. Refer to the 'Log Into the Command-Line Interface on the Device' section in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

View the cluster on **Devices** > **Device Management** to see which device is the control node.

Step 2 View how active and backup sessions are distributed across the cluster.

show cluster VPN-sessiondb distribution

Example:

Distribution information displays as follows:

```
> show cluster vpn-sessiondb distribution
Member 0 (unit-1-1): active: 209; backups at: 1(111), 2(98)
Member 1 (unit-1-3): active: 204; backups at: 0(108), 2(96)
Member 2 (unit-1-2): active: 0
```

Each row contains the member ID, member name, number of active sessions, and on which members the backup sessions reside. For the example above, one would read the information as:

- Member 0 has 209 active sessions, 111 sessions are backed up on member 1, 98 sessions are backed up on member 2
- Member 1 has 204 active sessions, 108 sessions are backed up on member 0, 96 sessions are backed up on member 2
- Member 2 has NO active sessions; therefore, NO cluster members are backing up sessions for this node. This member has recently joined the cluster.

Step 3 Redistribute sessions.

cluster redistribute VPN-sessiondb

Example:

```
> cluster redistribute vpn-sessiondb

Session redistribution initiated.
Use 'show cluster vpn-sessiondb distribution' to view distribution.
```

>

Depending on the number of sessions to redistribute and the load on the cluster, this may take some time. Syslogs containing the following phrases (and other system details not shown here) are provided as redistribution activity occurs:

Syslog Phrase	Notes
VPN session redistribution started	Control node only
Sent request to move *number* sessions from *orig-member-name* to *dest-member-name*	Control node only
Failed to send session redistribution message to *member-name*	Control node only
Received request to move *number* sessions from *orig-member-name* to *dest-member-name*	Data node only
Moved *number* sessions to *member-name*	The number of active sessions moved to the named cluster.
Failed to receive session move response from *dest-member-name*	Control node only
VPN session completed	Control node only
Cluster topology change detected. VPN session redistribution aborted.	

Step 4 Re-enter the **show cluster VPN-sessiondb distribution** command to view the results.

Cluster node management

After you deploy the cluster, you can change the configuration and manage cluster nodes.

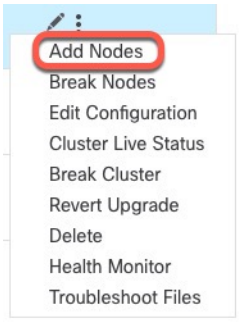
Add a new cluster node

You can add one or more new cluster nodes to an existing cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster, and choose **Add Nodes**.

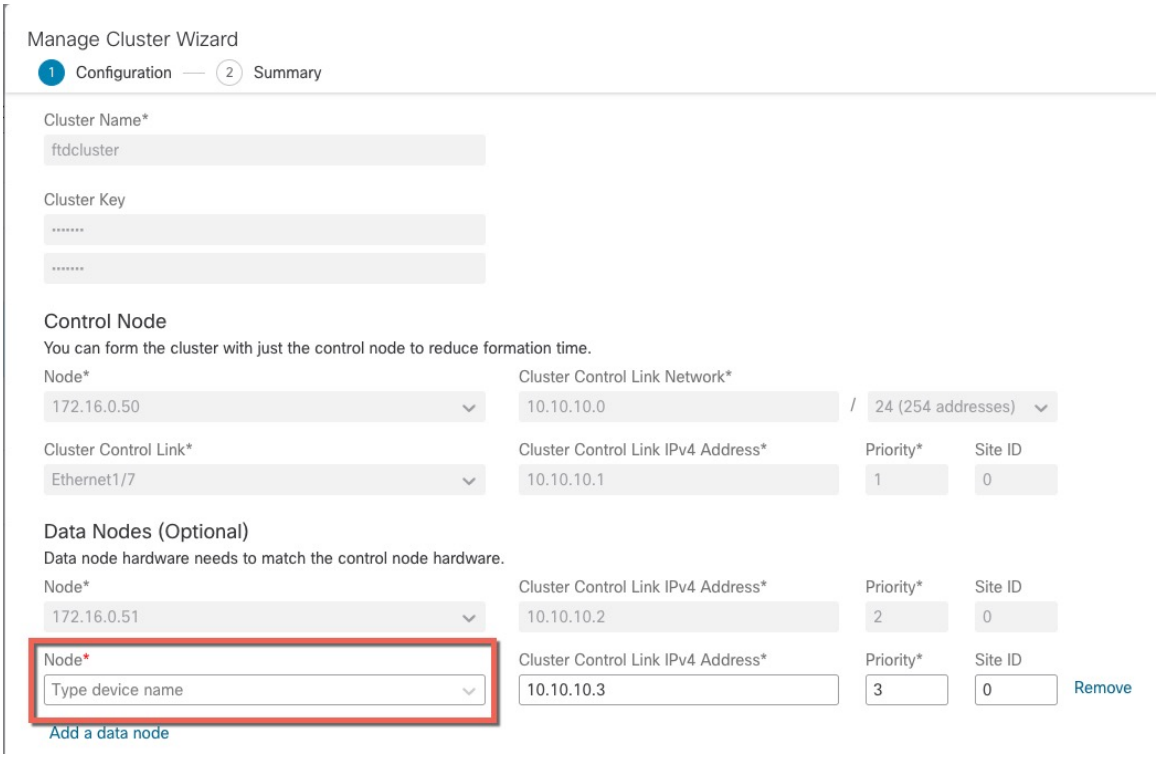
Figure 120: Add nodes



The **Manage Cluster Wizard** appears.

Step 2 From the **Node** menu, choose a device, adjust the IP address, priority, and Site ID if desired.

Figure 121: Manage cluster wizard

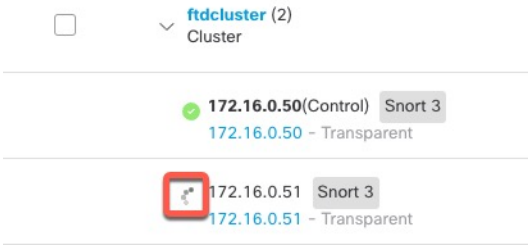


Step 3 To add additional nodes, click **Add a data node**.

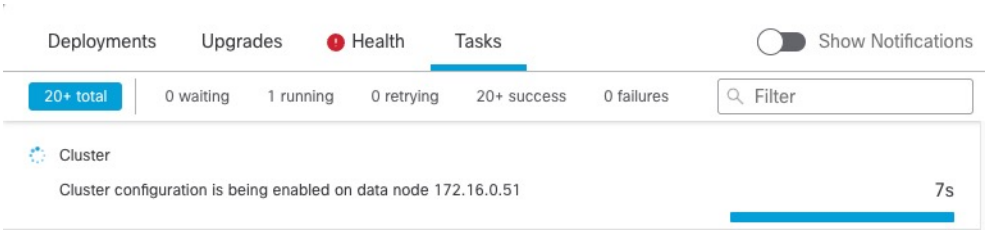
Step 4 Click **Continue**. Review the **Summary**, and then click **Save**

The node that is currently registering shows the loading icon.

Figure 122: Node registration



You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**.



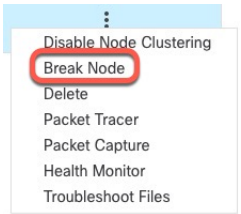
Break a node

You can remove a node from the cluster so that it becomes a standalone device. You cannot break the control node unless you break the entire cluster. The data node has its configuration erased.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the node you want to break, and choose **Break Node**.

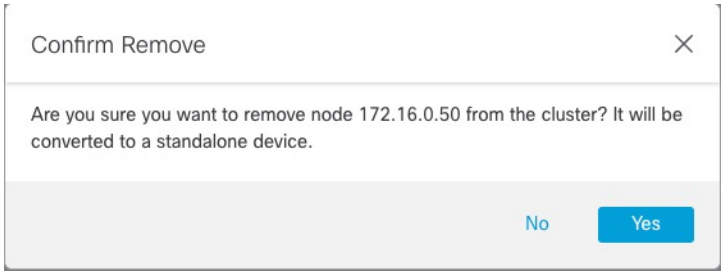
Figure 123: Break a node



You can optionally break one or more nodes from the cluster More menu by choosing **Break Nodes**.

Step 2 You are prompted to confirm the break; click **Yes**.

Figure 124: Confirm break



You can monitor the cluster node break by clicking the **Notifications** icon and choosing **Tasks**.

Break the cluster

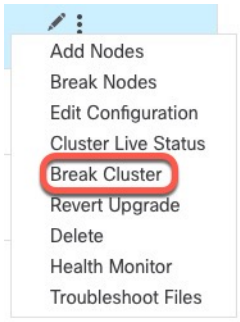
You can break the cluster and convert all nodes to standalone devices. The control node retains the interface and security policy configuration, while data nodes have their configuration erased.

Procedure

Step 1 Make sure all cluster nodes are being managed by the Cloud-Delivered Firewall Management Center by reconciling nodes. For more information, refer to Reconcile cluster nodes, on page 769.

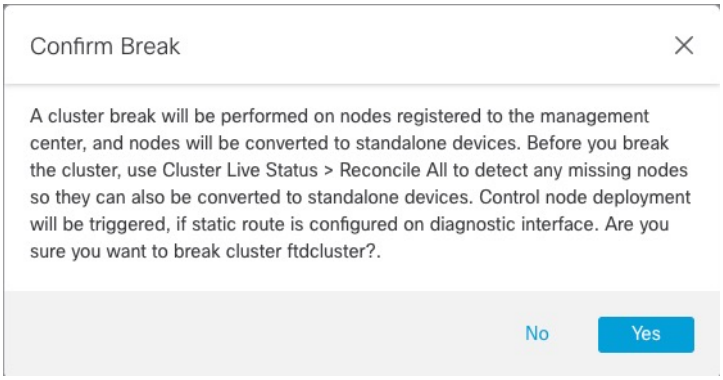
Step 2 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster, and choose **Break Cluster**.

Figure 125: Break cluster



Step 3 You are prompted to break the cluster; click **Yes**.

Figure 126: Confirm break



You can monitor the cluster break by clicking the **Notifications** icon and choosing **Tasks**.

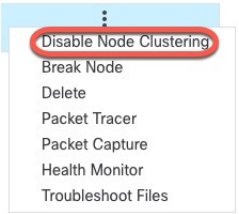
Disable clustering

You may want to deactivate a node in preparation for deleting the node, or temporarily for maintenance. This procedure is meant to temporarily deactivate a node; the node will still appear in the Cloud-Delivered Firewall Management Center device list. When a node becomes inactive, all data interfaces are shut down.

Procedure

Step 1 For the unit you want to disable, choose **Devices** > **Device Management**, click **More** (⋮), and choose **Disable Node Clustering**.

Figure 127: Disable clustering



If you disable clustering on the control node, one of the data nodes will become the new control node. Note that for centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node. You cannot disable clustering on the control node if it is the only node in the cluster.

Step 2 Confirm that you want to disable clustering on the node.

The node will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.

Step 3 To reenable clustering, refer Rejoin the cluster, on page 767.

Rejoin the cluster

If a node was removed from the cluster, for example for a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management**, click **More** (⋮), and choose **Enable Node Clustering**.

Step 2 Confirm that you want to enable clustering on the unit.

Change the control node

Use this procedure when you need to specify the exact unit you want to become the control node. For centralized features, forcing a control node change will drop all connections, requiring you to re-establish the connections on the new control node.



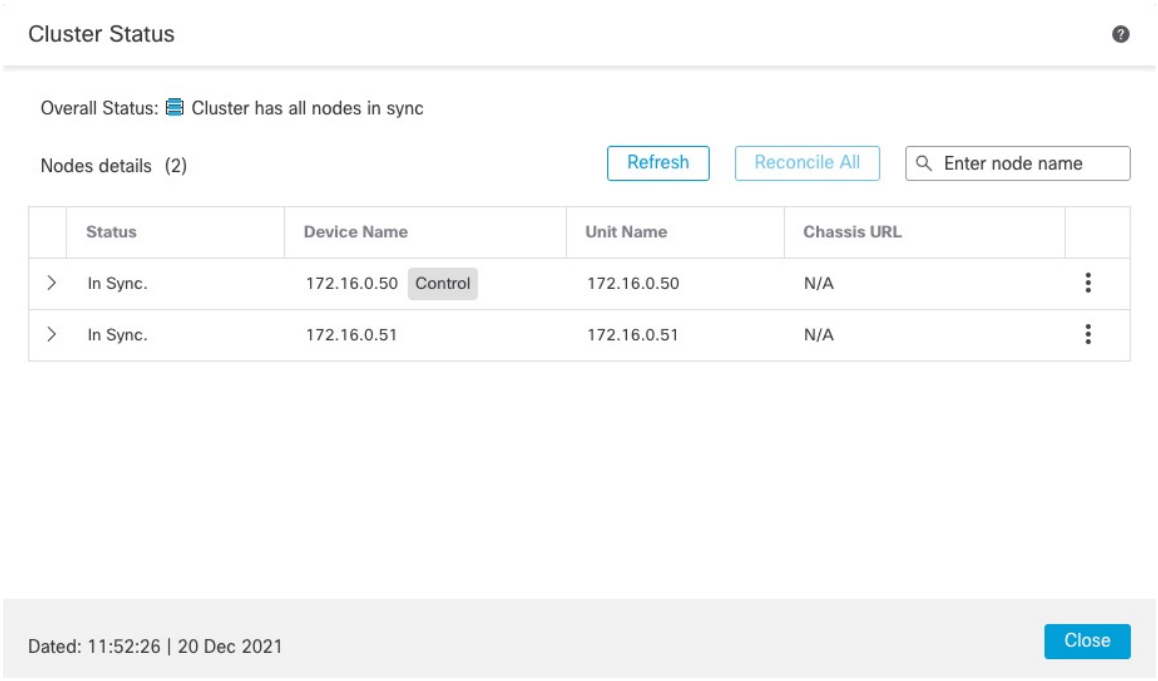
Caution The best method to change the control node is to disable clustering on the control node, wait for a new control election, and then re-enable clustering. If you must specify the *exact* unit you want to become the control node, use the procedure in this section. Note that for centralized features, if you force a control node change using either method, then all connections are dropped, and you have to re-establish the connections on the new control node.

Follow these steps to change the control node:

Procedure

Step 1 Open the **Cluster Status** dialog box by choosing **Devices** > **Device Management More** (⋮) **Cluster Live Status**.

Figure 128: Cluster status



Step 2 For the unit you want to become the control unit, choose **More** (⋮) **Change Role to Control**.

Step 3 You are prompted to confirm the role change. Check the check box, and click **OK**.

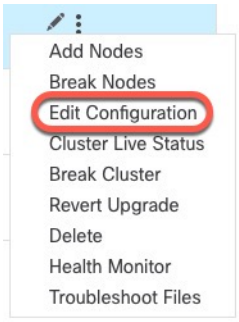
Edit the cluster configuration

You can edit the cluster configuration. If you change the cluster key, cluster control link interface, or cluster control link network, the cluster will be broken and reformed automatically. Until the cluster is reformed, you may experience traffic disruption. If you change the cluster control link IP address for a node, node priority, or site ID, only the affected nodes are broken and readded to the cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster, and choose **Edit Configuration**.

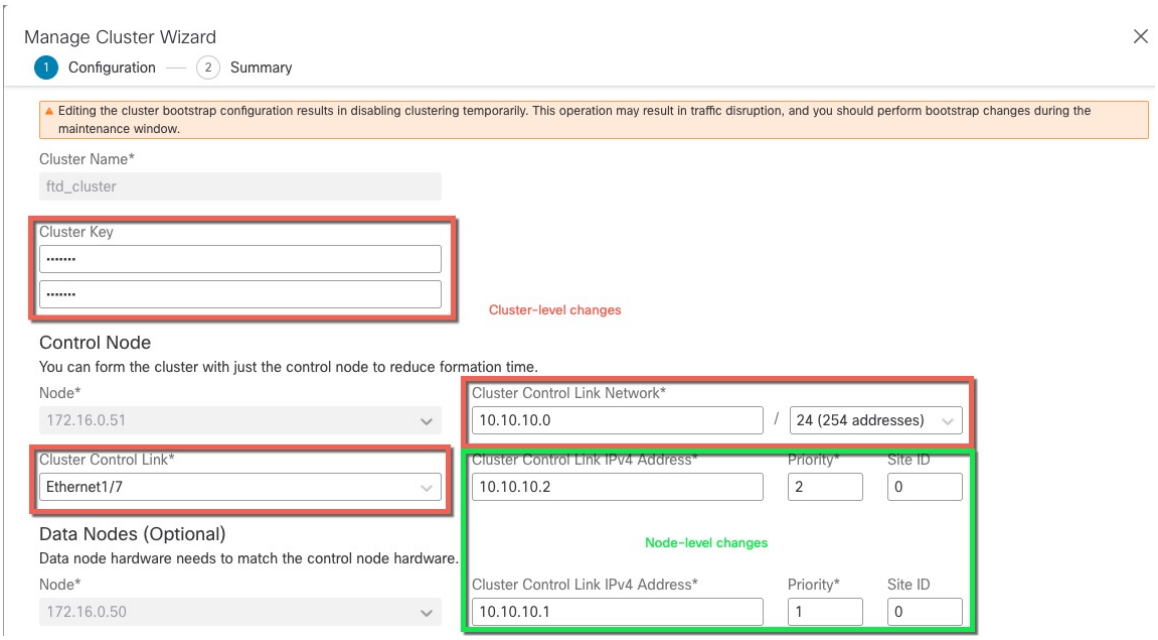
Figure 129: Edit configuration



The **Manage Cluster Wizard** appears.

Step 2 Update the cluster configuration.

Figure 130: Manage cluster wizard



If the cluster control link is an EtherChannel, you can edit the interface membership and LACP configuration by clicking **Edit** (✎) next to the interface drop-down menu.

Step 3 Click **Continue**. Review the **Summary**, and then click **Save**

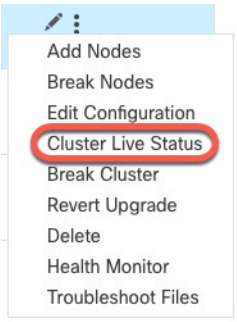
Reconcile cluster nodes

If a cluster node fails to register, you can reconcile the cluster membership from the device to the Cloud-Delivered Firewall Management Center. For example, a data node might fail to register if the Cloud-Delivered Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

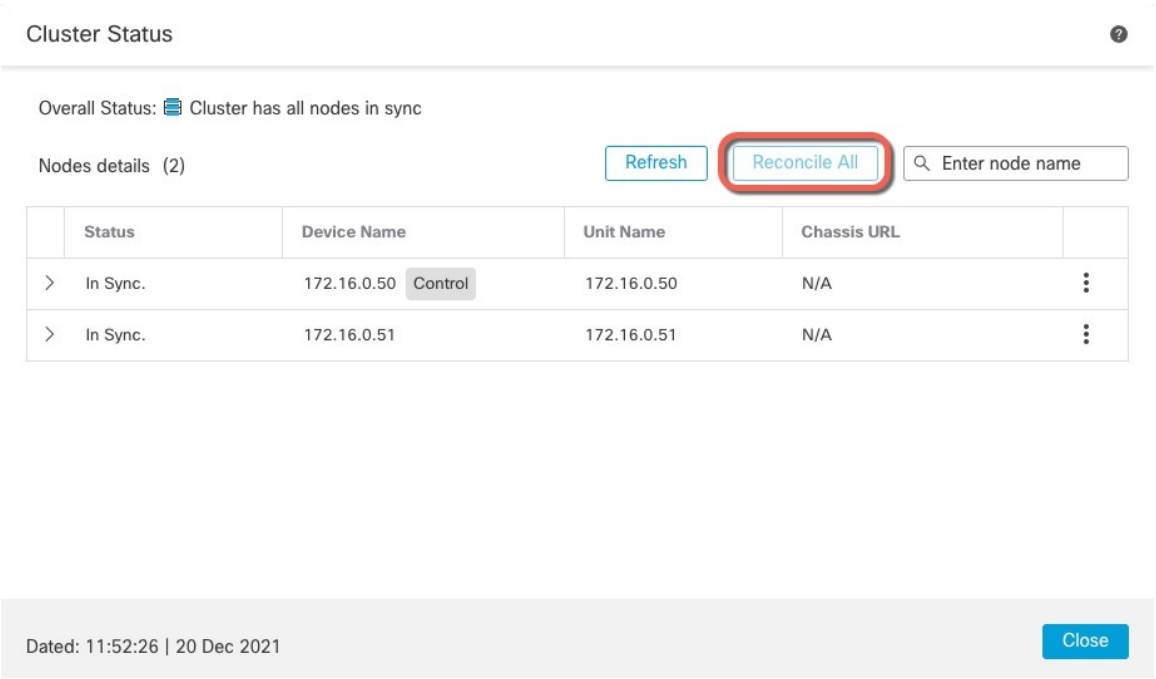
Step 1 Choose **Devices** > **Device Management More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Figure 131: Cluster live status



Step 2 Click **Reconcile All**.

Figure 132: Reconcile all



For more information about the cluster status, refer to Monitor the cluster, on page 772.

Unregister the Cluster or Nodes and Register to a New Cloud-Delivered Firewall Management Center

You can unregister the cluster from the Cloud-Delivered Firewall Management Center, which keeps the cluster intact. You might want to unregister the cluster if you want to add the cluster to a new Cloud-Delivered Firewall Management Center.

You can also unregister a node from the Cloud-Delivered Firewall Management Center without breaking the node from the cluster. Although the node is not visible in the Cloud-Delivered Firewall Management Center, it is still part of the cluster, and it will continue to pass traffic and could even become the control node. You cannot unregister the current control node. You might want to unregister the node if it is no longer reachable from the Cloud-Delivered Firewall Management Center, but you still want to keep it as part of the cluster while you troubleshoot management connectivity.

Unregistering a cluster:

- Severs all communication between the Cloud-Delivered Firewall Management Center and the cluster.
- Removes the cluster from the **Device Management** page.
- Returns the cluster to local time management if the cluster's platform settings policy is configured to receive time from the Cloud-Delivered Firewall Management Center using NTP.
- Leaves the configuration intact, so the cluster continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the cluster again to the same or a different Cloud-Delivered Firewall Management Center causes the configuration to be removed, so the cluster will stop processing traffic at that point; the cluster configuration remains intact so you can add the cluster as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.

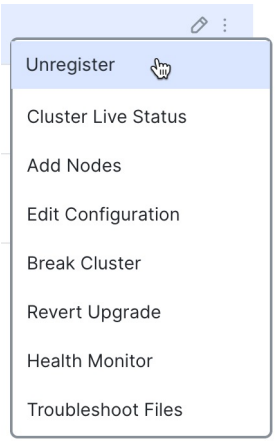
Before you begin

This procedure requires CLI access to one of the nodes.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster or node, and choose **Unregister**.

Figure 133: Unregister cluster or node



Step 2 You are prompted to unregister the cluster or node; click **Yes**.

Step 3 You can register the cluster to a new (or the same) Cloud-Delivered Firewall Management Center by adding one of the cluster members as a new device.

You only need to add one of the cluster nodes as a device, and the rest of the cluster nodes will be discovered.

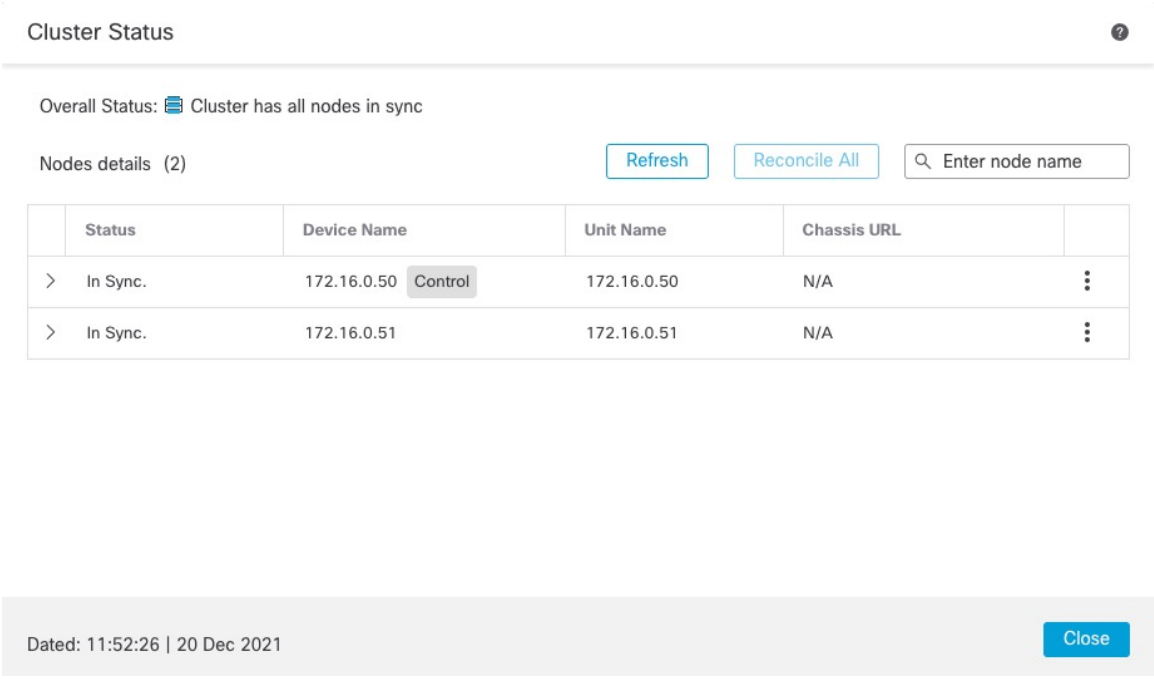
Step 4 To re-add an unregistered node, refer to Reconcile cluster nodes, on page 769.

Monitor the cluster

You can monitor the cluster in the Cloud-Delivered Firewall Management Center and at the Firewall Threat Defense CLI.

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management More** (⋮) icon or from the **Devices** > **Device Management** > **Cluster** page **General** area **Cluster Live Status** link.

Figure 134: Cluster status



The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes these states:

- **In Sync**—The node is registered with the Cloud-Delivered Firewall Management Center.
- **Pending Registration**—The node is part of the cluster, but has not yet registered with the Cloud-Delivered Firewall Management Center. If a node fails to register, you can retry registration by clicking **Reconcile All**.
- **Clustering is disabled**—The node is registered with the Cloud-Delivered Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the node from the cluster.
- **Joining cluster**—The node is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Cloud-Delivered Firewall Management Center.

For each node, you can view the **Summary** or the **History**.

Figure 135: Node summary

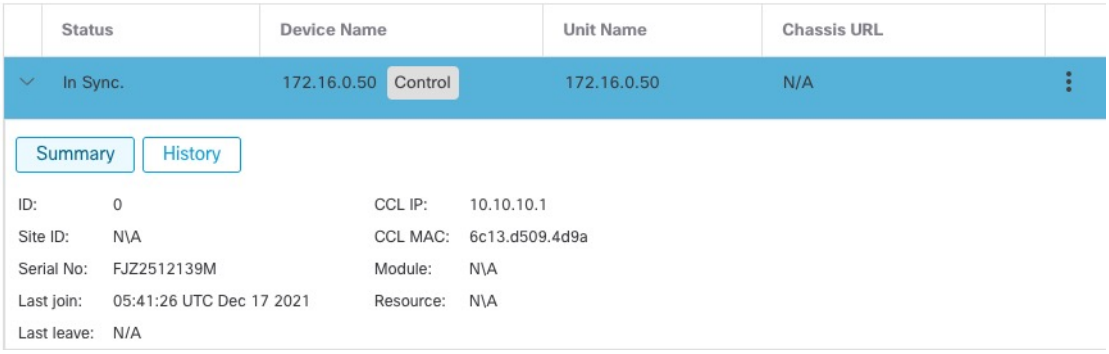
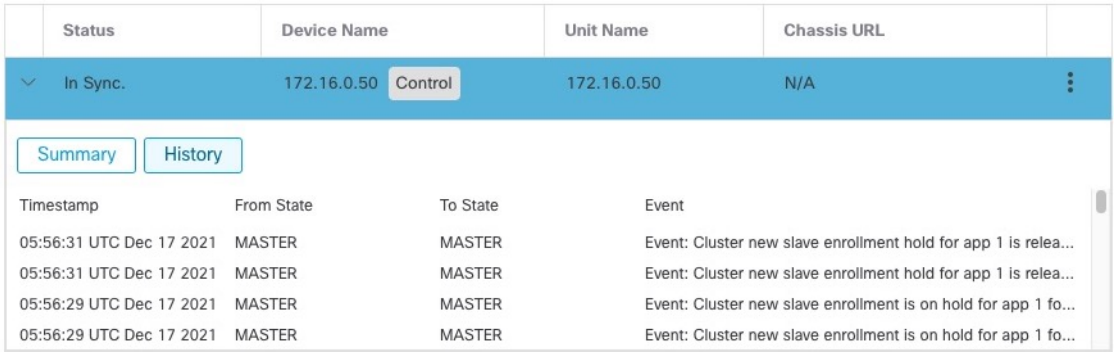


Figure 136: Node history



- **System** (⚙) > **Tasks** page.

  The **Tasks** page shows updates of the Cluster Registration task as each node registers.

- **Devices** > **Device Management***cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member nodes, including the control node shown with its role next to the IP address. For nodes that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Cluster health monitor dashboard

The cluster health monitor dashboard provides visibility into cluster status, metrics, and statistics.

View cluster health

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

- You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.
- Ensure you have created a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Procedure

	Command or Action	Purpose
Step 1	Choose **System** (⚙) > **Health** > **Monitor**.	Use the Monitoring navigation pane to access node-specific health monitors.
Step 2	In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed cluster devices.	
Step 3	To view the cluster health statistics, click on the cluster name. The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:	• Overview— Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information. • Load Distribution— Traffic and packet distribution across the cluster nodes. • Member Performance— Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation. • CCL— Interface status and aggregate traffic statistics. You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.
Step 4	You can configure the time range from the drop-down in the upper-right corner. The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.	Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

	Command or Action	Purpose
Step 5	Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.	The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.
Step 6	(For node-specific health monitor) View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.	Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.
Step 7	(For node-specific health monitor) The device monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:	• Overview — Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information. • CPU — CPU utilization, including the CPU usage by process and by physical cores. • Memory — Device memory utilization, including data plane and Snort memory usage. • Interfaces — Interface status and aggregate traffic statistics. • Connections — Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts. • Snort — Statistics that are related to the Snort process. • ASP drops — Statistics related to the dropped packets against various reasons. You can navigate through the various metrics dashboards by clicking on the labels. For more information, refer to Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.
Step 8	Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your	For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

	Command or Action	Purpose
	own variable set from the available metric groups.	

Cluster metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 52: Cluster metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number
Packets	Packet distribution count in the cluster for every second.	number

Troubleshoot the cluster

Use these methods to troubleshoot the devices and clusters:

- Use **Logs** to generates troubleshooting files for all nodes in the cluster. Use this option when you need one cluster-level collection instead of generating troubleshooting files from each node separately.
- Use **CLI** to run a predefined set of cluster diagnostic commands and displays the output. Use this option to check cluster configuration, member state, cluster health, control-plane transport statistics, packet drops, interface status, CPU usage, memory buffers, ARP entries, and reachability. The **CLI** option displays command output; it does not open an interactive CLI session.

Commands
show running-config cluster: Review the cluster group configuration, local unit, cluster interface, priority, health-check, auto-rejoin, and cLACP settings.

Commands
show cluster info: Confirm that clustering is enabled and identify the control node, member state, software version, serial number, CCL IP address, CCL MAC address, and last join or leave state.
show cluster info health: Check unit, interface, service, and overall cluster health.
show cluster info transport cp: Review control-plane transport counters, including reliable and unreliable message counters, errors, retransmits, drops, and buffer conditions.
show version: Capture platform, software, build, UUID, LSP, VDB, uptime, resource, and license context.
show asp drop: Identify accelerated security path packet or flow drop reasons and counters.
show counters: Review protocol, feature, and cluster CCL counters.
show arp: Verify ARP entries for interfaces that participate in cluster communication.
show int ip brief: Check interface IP addresses and administrative and line protocol status.
show blocks: Check buffer availability and failed allocations.
show cpu detailed: Check data path, control point, and total CPU usage.
show interface *ccl_interface*: Review interface state, MTU, MAC address, IP address, traffic counters, errors, and drops for the cluster-related interface.
ping *ccl_ip* **size** *ccl_mtu* **repeat 2**: Test reachability to a cluster peer by using the packet size selected by the tool.

- Use the **CCL Ping** when you need to test cluster control link connectivity between joined cluster nodes.
- Troubleshooting files—If a node fails to join the cluster, a troubleshooting file is automatically generated. You can also generate and download troubleshooting files from the **Devices** > **Device Management** and then choose **Add**, **ClusterGeneral** area. For more information, refer to Generate troubleshooting files.

  You can also generate files from the **Device Management** page. Click **More** (⋮) and choose **Troubleshoot Files**.
- CLI output—Choose **Devices** > **Device Management** and then choose **Add**, **ClusterGeneral** area, you can view a set of pre-defined CLI outputs that can help you troubleshoot the cluster. These commands are automatically run for the cluster:

  You can also enter any **show** command in the **Command** field.

Perform a ping on the cluster control link

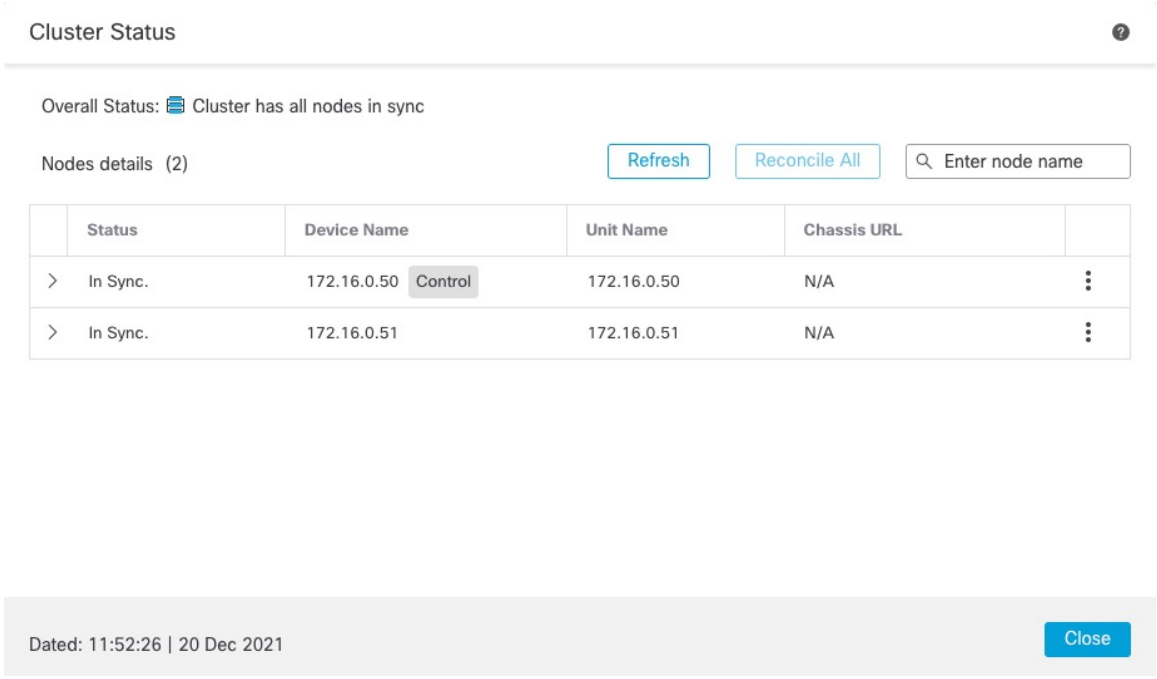
When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the initial ping fails, the node tries a ping using a smaller packet size (the MTU divided by 2, then by 4, then by 8) until a ping succeeds. A notification is generated so you can fix the MTU mismatch on connecting switches and try again. This tool lets you manually ping all nodes that have already joined the cluster in case you are having cluster control link connectivity problems.

You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs.

Procedure

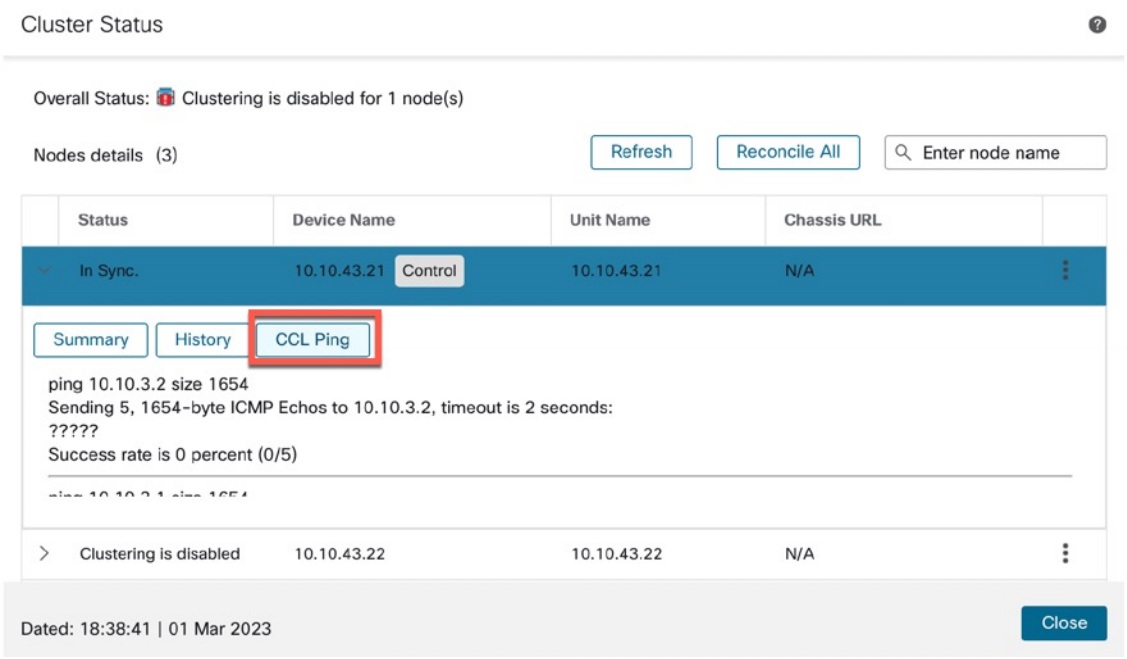
Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) icon next to the cluster, and choose **Cluster Live Status**.

Figure 137: Cluster status



Step 2 Expand one of the nodes, and click **CCL Ping**.

Figure 138: CCL ping



The node sends a ping on the cluster control link to every other node using a packet size that matches the maximum MTU.

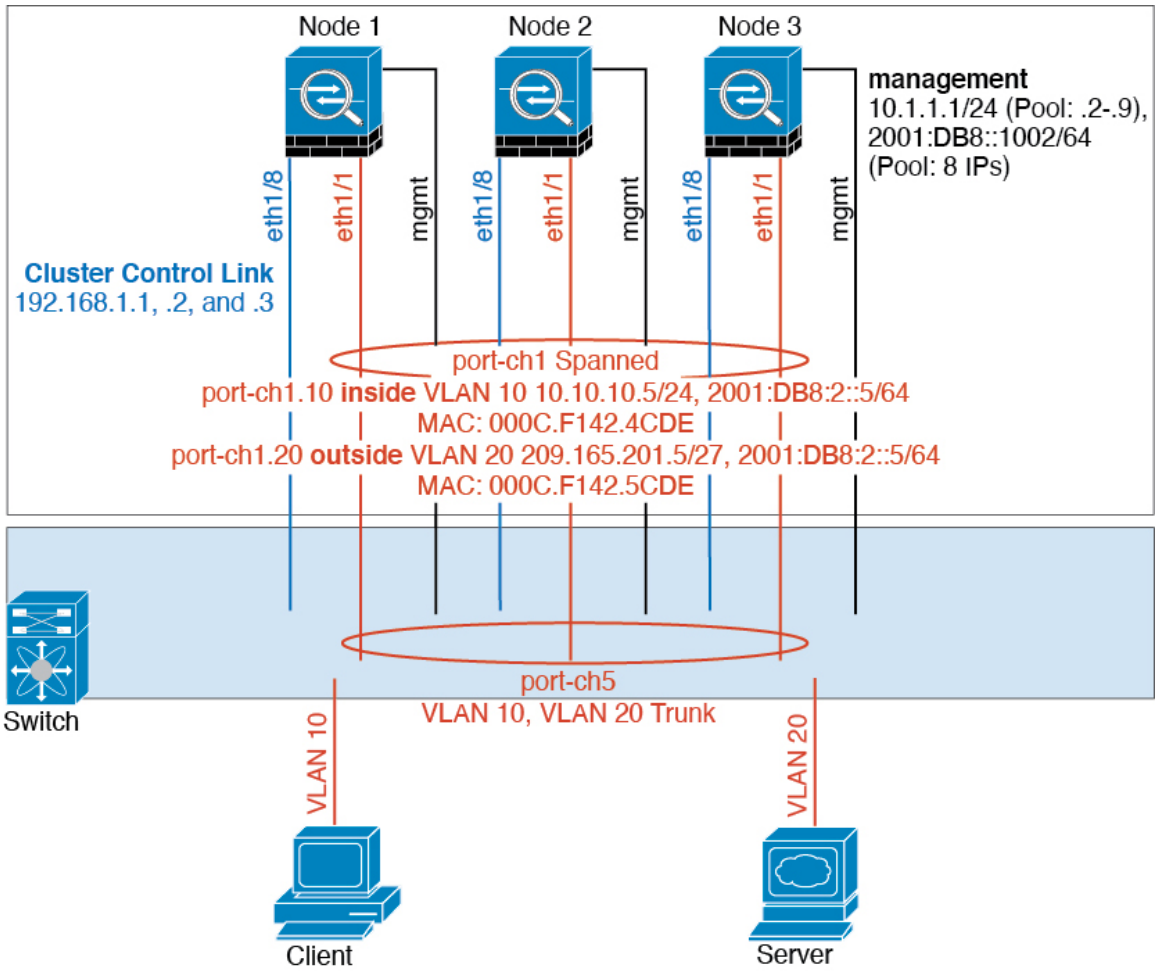
Clustering examples

These examples include examples for typical deployments.

Firewall on a stick

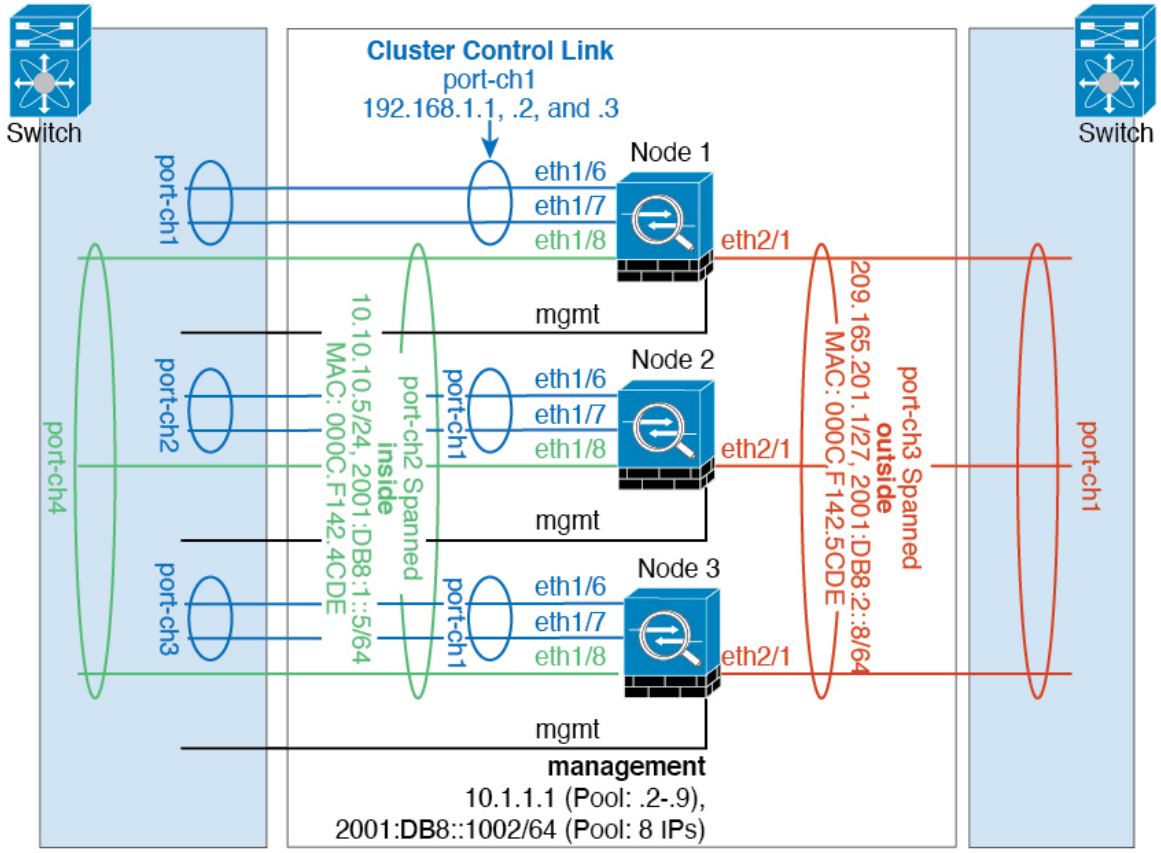
Data traffic from different security domains are associated with different VLANs. For example, VLAN 10 for the inside network and VLAN 20 for the outside network. Each Threat Defense has a single physical port connected to the external switch or router. Trunking is enabled so that all packets on the physical link are 802.1q encapsulated. The Threat Defense is the firewall between VLAN 10 and VLAN 20.

When using Spanned EtherChannels, all data links are grouped into one EtherChannel on the switch side. If the Threat Defense becomes unavailable, the switch will rebalance traffic between the remaining units.



Traffic segregation

You may prefer physical separation of traffic between the inside and outside network.



As shown in the diagram above, there is one Spanned EtherChannel on the left side that connects to the inside switch, and the other on the right side to outside switch. You can also create VLAN subinterfaces on each EtherChannel if desired.

Reference for Clustering

This section includes more information about how clustering operates.

Firewall Threat Defense features and clustering

Some Firewall Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported features with clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, refer to the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Cloud-Delivered Firewall Management Center GUI. For more information, refer to FlexConfig policies, on page 1309.

- Remote access VPN (SSL VPN and IPsec VPN)
- DHCP client, server, and proxy. DHCP relay is supported.
- Virtual Tunnel Interfaces (VTIs)
- High Availability
- Integrated Routing and Bridging
- Cloud-Delivered Firewall Management Center UCAPL/CC mode
- DHCP client, server, and proxy. DHCP relay is supported.

Centralized features for clustering

These features are only supported on the control node, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member nodes to the control node over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control nodes before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control node.

For centralized features, if the control node fails, all connections are dropped, and you have to re-establish the connections on the new control node.

Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, refer to the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Cloud-Delivered Firewall Management Center GUI. For more information, refer to FlexConfig policies, on page 1309.

- These application inspections are supported:
  - DCERPC and ESMTP
  - NetBIOS and PPTP
  - RSH
  - SQLNET and SUNRPC
  - TFTP

- XDMCP

- Static route monitoring
- Site-to-site VPN
- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Dynamic routing

Connection settings and clustering

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

FTP and Clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

Multicast routing in individual interface mode

In Individual interface mode, units do not act independently with multicast. All data and routing packets are processed and forwarded by the control unit, thus avoiding packet replication.

NAT and clustering

NAT can impact the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different devices in the cluster because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the device that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link.

Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

Guidelines for NAT and clustering

General guidelines for NAT and clustering

- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number

of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.

- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

PAT with port block allocation guidelines

- Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
- Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
- On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
- When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.

Limitations for NAT and Clustering

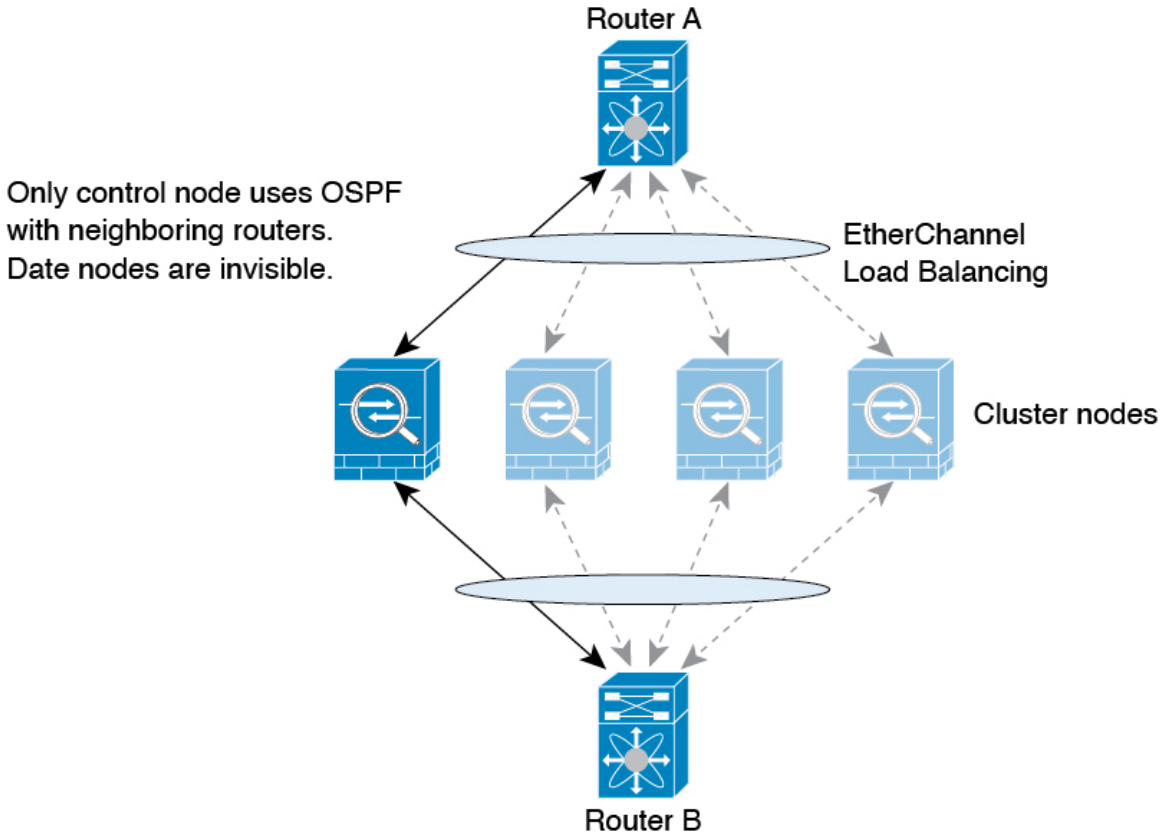
- No static PAT for these inspections:
  - FTP
  - RSH
  - SQLNET
  - TFTP
  - XDMCP
  - SIP

- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.
- No Proxy ARP—For Individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with a node that may no longer be in the cluster. The upstream router needs a static route or PBR with Object Tracking for the mapped addresses that points to the Main cluster IP address. This is not an issue for a Spanned EtherChannel, because there is only one IP address associated with the cluster interface.
- No interface PAT on an Individual interface—Interface PAT is not supported for Individual interfaces.

Dynamic routing

The routing process only runs on the control node, and routes are learned through the control node and replicated to data nodes. If a routing packet arrives at a data node, it is redirected to the control node.

Figure 139: Dynamic routing in spanned EtherChannel mode



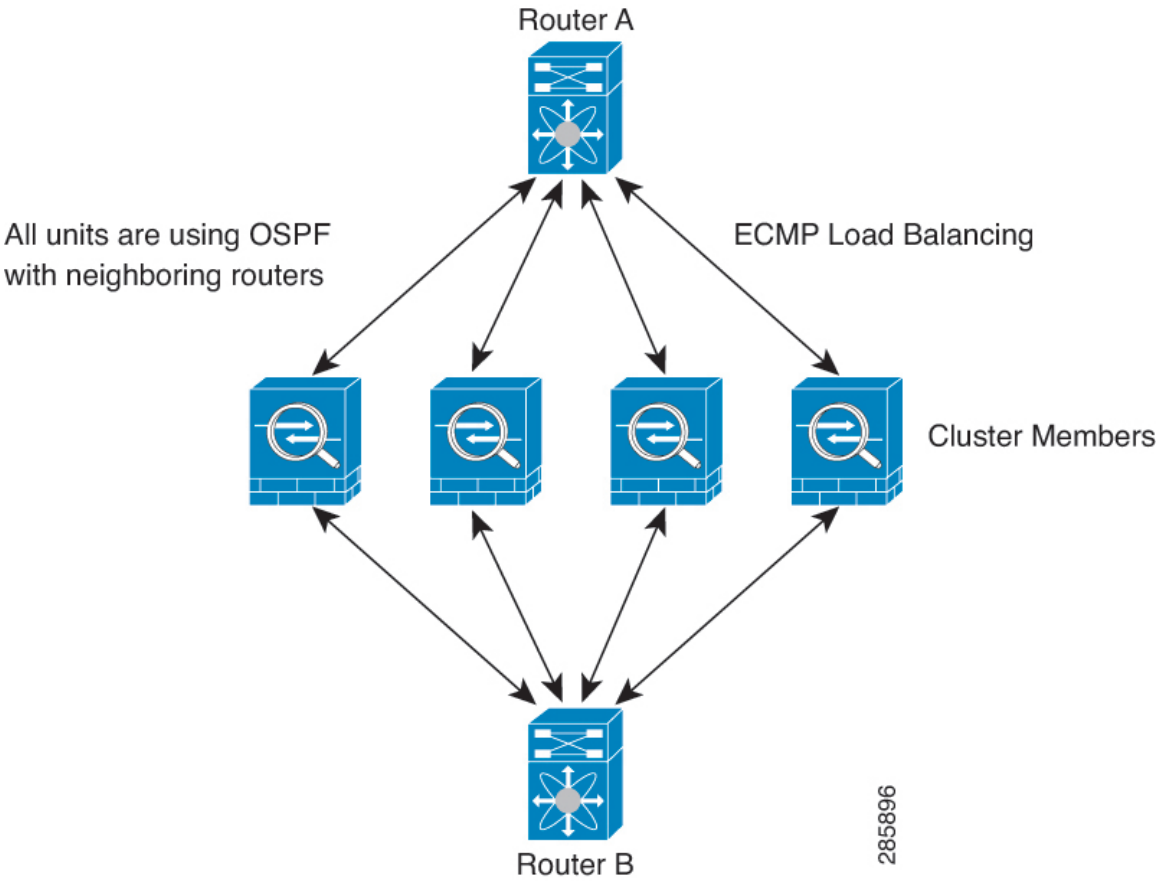
After the data node learn the routes from the control node, each node makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the control node to data nodes. If there is a control node switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

Dynamic routing in individual interface mode

In Individual interface mode, each node runs the routing protocol as a standalone router, and routes are learned by each node independently.

Figure 140: Dynamic routing in individual interface mode



In the above diagram, Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when talking to external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

EIGRP does not form neighbor relationships with cluster peers in individual interface mode.

Note If the cluster has multiple adjacencies to the same router for redundancy purposes, asymmetric routing can lead to unacceptable traffic loss. To avoid asymmetric routing, group all of these node interfaces into the same traffic zone. For more information, refer to Create an ECMP zone, on page 1433.

SIP inspection and clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and clustering

You should always use the local address, and not the main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and clustering

Syslog and clustering is a network management configuration that enables logging functionality within clustered network environments.

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

Cisco TrustSec and clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

VPN and clustering

Site-to-site VPN is a centralized feature; only the control node supports VPN connections.

Note Remote access VPN is not supported with clustering.

VPN functionality is limited to the control node and does not take advantage of the cluster high availability capabilities. If the control node fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control node is elected, you must reestablish the VPN connections.

When you connect a VPN tunnel to a Spanned EtherChannel address, connections are automatically forwarded to the control node. For connections to an Individual interface when using PBR or ECMP, you must always connect to the Main cluster IP address, not a Local address.

VPN-related keys and certificates are replicated to all nodes.

Performance scaling factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

Throughput calculation

For example, if your model can handle approximately 10 Gbps of traffic when running alone, then for a cluster of 8 units, the maximum combined throughput will be approximately 80% of 80 Gbps (8 units x 10 Gbps): 64 Gbps.

Control node election

Summary

The key components involved in the control node election process are:

- Cluster nodes: Communicate over the cluster control link and participate in the election process
- Election requests: Broadcast every 3 seconds by nodes when clustering is enabled
- Priority values: Set between 1 and 100 where 1 is the highest priority
- Control node: The elected node that coordinates cluster operations

Workflow

These stages describe how cluster nodes elect a control node:

1. When clustering is enabled for a node (or when it first starts up with clustering already enabled), the node broadcasts an election request every 3 seconds.
2. Any other nodes with a higher priority respond to the election request. The priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a node does not receive a response from another node with a higher priority, then it becomes the control node.



Note If multiple nodes tie for the highest priority, the cluster node name and then the serial number is used to determine the control node.

4. If a node later joins the cluster with a higher priority, it does not automatically become the control node. The existing control node always remains as the control node unless it stops responding, at which point a new control node is elected.
5. In a "split brain" scenario when there are temporarily multiple control nodes, the node with highest priority retains the role while the other nodes return to data node roles.

Note You can manually force a node to become the control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node.

High availability within clusters

Clustering provides high availability by monitoring node and interface health and by replicating connection states between nodes.

Node health monitoring

Each node periodically sends a broadcast heartbeat packet over the cluster control link. If the control node does not receive any heartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining nodes.

Split brain scenario

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role.

Interface monitoring

Each node monitors the link status of all named hardware interfaces in use, and reports status changes to the control node.

- **Spanned EtherChannel**—Uses cluster Link Aggregation Control Protocol (cLACP). Each node monitors the link status and the cLACP protocol messages to determine if the port is still active in the EtherChannel. The status is reported to the control node.
- **Individual interfaces** (Routed mode only)—Each node self-monitors its interfaces and reports interface status to the control node.

When you enable health monitoring, all physical interfaces (including the main EtherChannel) are monitored by default; you can optionally disable monitoring per interface. Only named interfaces can be monitored. For example, the named EtherChannel must fail to be considered failed, which means all member ports of an EtherChannel must fail to trigger cluster removal.

A node is removed from the cluster if its monitored interfaces fail. The amount of time before the Firewall Threat Defense removes a member from the cluster depends on the type of interface and whether the node is an established member or is joining the cluster. For EtherChannels (spanned or not): If the interface is down on an established member, then the Firewall Threat Defense removes the member after 9 seconds. The Firewall Threat Defense does not monitor interfaces for the first 90 seconds that a node joins the cluster. Interface

status changes during this time will not cause the Firewall Threat Defense to be removed from the cluster. For non-EtherChannels, the node is removed after 500 ms, regardless of the member state.

Status after failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The device automatically tries to rejoin the cluster, depending on the failure event.



Note When the device becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoin the cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firewall Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firewall Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firewall Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from Cloud-Delivered Firewall Management Center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.

Data path connection state replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the

connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 53: Features replicated across the cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	No	—

How the cluster manages connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection roles

These roles are defined for each connection:

- **Owner**—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- **Backup owner**—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- **Director**—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the

owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

ICMP/ICMPv6 hash details:

- For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
- For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
- For other packets, both source and destination ports are 0.

- **Forwarder**—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.



Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- **Fragment Owner**—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

New connection ownership

When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. If a reverse flow arrives at a different node, it is redirected back to the original node.

Traffic redirection is not supported in this release. When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. All the subsequent packets for the same connection should arrive the same node. If any connection packets arrive at a different node, they will be

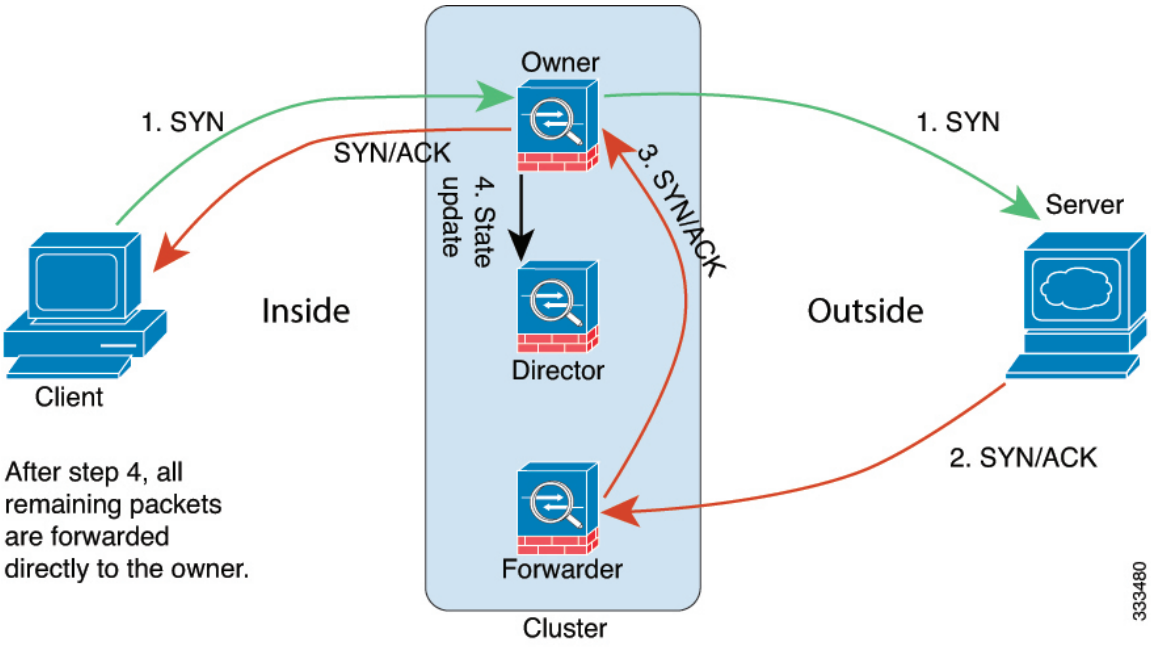
dropped. If a reverse flow arrives at a different node, it will be dropped as well. For centralized features, if the connections do not arrive on the control node, they will be dropped.

Sample data flow for TCP

Summary

This example shows the establishment of a new connection.

Workflow



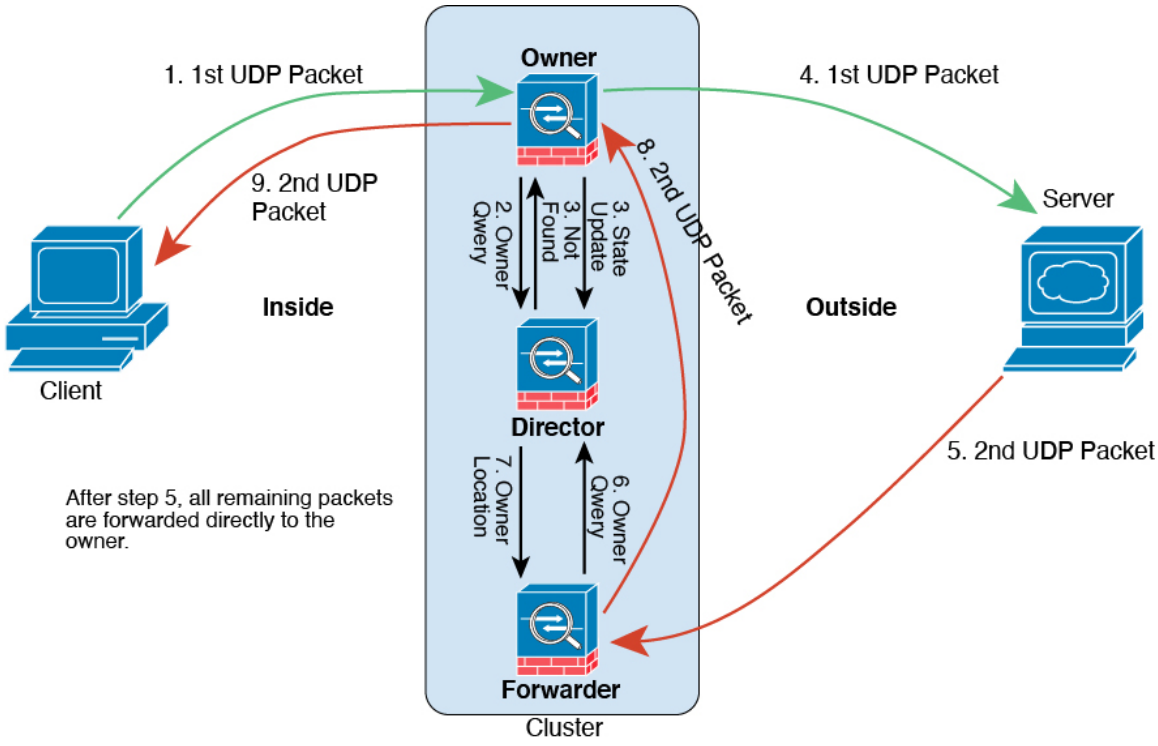
1. The SYN packet originates from the client and is delivered to one device (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different device (based on the load balancing method). This device is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample data flows for ICMP and UDP

Summary

This example shows the establishment of a new connection.

Workflow



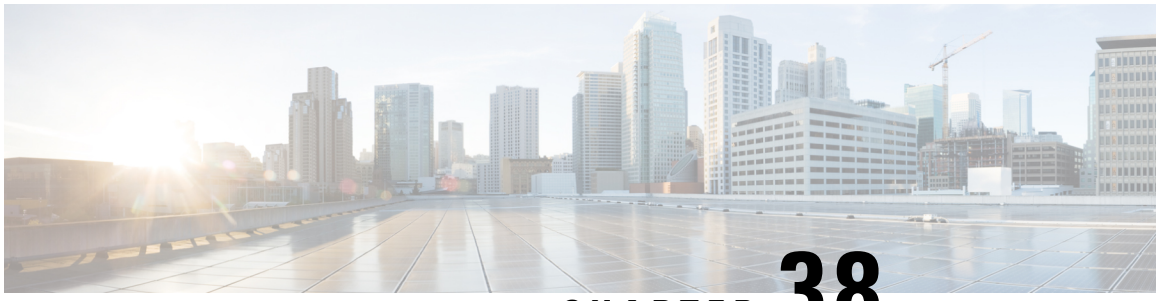
1. The first UDP packet originates from the client and is delivered to one device (based on the load balancing method).
2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for Clustering

Lists the clustering features introduced in each release.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
16-node clusters for the Secure Firewall 3100/4200.	20241030	7.6.0	For the Secure Firewall 3100 and 4200, the maximum nodes were increased from 8 to 16. See: Clustering for the Secure Firewall 3100/4200
Individual interface mode for Secure Firewall 3100/4200 clusters.	20241030	7.6.0	Individual interfaces are normal routed interfaces, each with their own local IP address used for routing. The main cluster IP address for each interface is a fixed address that always belongs to the control node. When the control node changes, the main cluster IP address moves to the new control node, so management of the cluster continues seamlessly. Load balancing must be configured separately on the upstream switch. Restrictions: Not supported for container instances. New/modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **Cluster** > **Interfaces / EIGRP / OSPF / OSPFv3 / BGP** • **Objects** > **Object Management** > **Address Pools** > **MAC Address Pool** See: Clustering for the Secure Firewall 3100/4200 and Address Pools
MTU ping test from data node on node join	20241030	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size twice the cluster control link MTU. Previously, only the control node sent a ping. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again. Added/modified commands: **show cluster history**. See: Clustering for the Secure Firewall 3100/4200, Clustering for Threat Defense Virtual in a Private Cloud, Clustering for Threat Defense Virtual in a Public Cloud, Clustering for the Firepower 4100/9300
Clustering for the Secure Firewall 4200	20230929	7.4.0	The Secure Firewall 4200 supports Spanned EtherChannel clustering for up to 8 nodes.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Cluster health monitor settings	20221213	Any	You can now edit cluster health monitor settings. New/modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard	20221213	Any	You can now view cluster health on the cluster health monitor dashboard. New/modified screens: **System** > **Health** > **Monitor**
Automatic configuration of the cluster control link MTU	20220609	7.2.0	The MTU of the cluster control link interface is now automatically set to 100 bytes more than the highest data interface MTU; by default, the MTU is 1600 bytes.
Clustering for the Secure Firewall 3100	20220609	7.1.0	The Secure Firewall 3100 supports Spanned EtherChannel clustering for up to 8 nodes. New/modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Secure Firewall 3100



CHAPTER 38

Clustering: Private Cloud

Clustering lets you group multiple Firewall Threat Defense Virtuals together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. You can deploy Firewall Threat Defense Virtual clusters in a private cloud using VMware and KVM. Only routed firewall mode is supported.

Note Some features are not supported when using clustering. See Unsupported features and clustering, on page 832.

- About clustering: private cloud, on page 799
- Licenses for clustering: Private cloud, on page 803
- Prerequisites for clustering: Private cloud, on page 804
- Guidelines for clustering: Private cloud, on page 805
- Configure clustering: private cloud, on page 806
- Manage cluster nodes, on page 817
- Monitor cluster status, on page 826
- Reference for clustering, on page 832
- History for clustering: private cloud, on page 844

About clustering: private cloud

This section describes the clustering architecture and how it works.

How the cluster fits into your network

The cluster consists of multiple firewalls acting as a single device. To act as a cluster, the firewalls need the following infrastructure:

- Isolated network for intra-cluster communication, known as the cluster control link, using VXLAN interfaces. VXLANs, which act as Layer 2 virtual networks over Layer 3 physical networks, let the Firewall Threat Defense Virtual send broadcast/multicast messages over the cluster control link.
- Management access to each firewall for configuration and monitoring. The Firewall Threat Defense Virtual deployment includes a Management 0/0 interface that you will use to manage the cluster nodes.

When you place the cluster in your network, the upstream and downstream routers need to be able to load-balance the data coming to and from the cluster using Layer 3 Individual interfaces and one of the following methods:

- Policy-Based Routing—The upstream and downstream routers perform load balancing between nodes using route maps and ACLs.
- Equal-Cost Multi-Path Routing—The upstream and downstream routers perform load balancing between nodes using equal cost static or dynamic routes.

Layer 2 Spanned EtherChannels are not supported.

Control and data node roles

When you add a cluster in the Cloud-Delivered Firewall Management Center , you choose one firewall to be the control node, and all additional firewalls to be data nodes. When you first create the cluster, the control node you specify will become the control node simply because it is the first node added to the cluster. Later, if multiple cluster nodes come online at the same time, the control node is determined by the priority setting; the priority is set between 1 and 100, where 1 is the highest priority.

When you create the cluster and deploy the cluster configuration, the Cloud-Delivered Firewall Management Center deploys a special bootstrap configuration to each node that includes fundamental settings such as the cluster control link configuration. The bootstrap configuration lets each node join the cluster. Most bootstrap settings are defined in the cluster wizard; however, you define the cluster control link interface hardware configuration (for example, creating an EtherChannel or setting the Ethernet speed) on the control node before you create the cluster. The cluster control link settings are copied to the bootstrap configuration for each node.

All nodes in the cluster share the same policy configuration. The node that you initially specify as the control node will overwrite the configuration on the data nodes when they join the cluster, so you only need to perform initial configuration on the control node before you form the cluster. After you create the cluster, all configuration changes are made at the cluster level and are shared by all nodes.

Some features do not scale in a cluster, and the control node handles all traffic for those features.

Individual interfaces

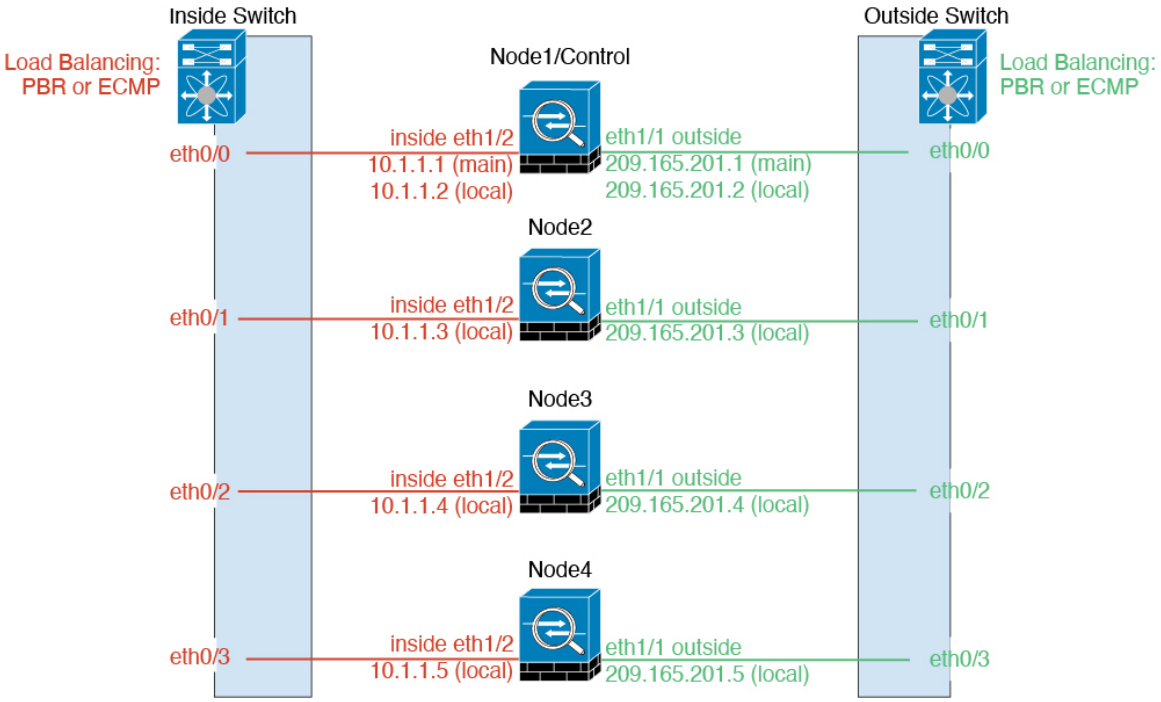
You can configure cluster interfaces as individual interfaces.

Individual interfaces are normal routed interfaces, each with their own *local IP address* used for routing. The *main cluster IP address* for each interface is a fixed address that always belongs to the control node. When the control node changes, the main cluster IP address moves to the new control node, so management of the cluster continues seamlessly.

IPS-only interfaces (inline sets and passive interfaces) are not supported as Individual interfaces.

Because interface configuration must be done only on the control node, configure a pool of IP addresses for a given interface on the cluster nodes, including one for the control node.

Load balancing must be configured separately on the upstream switch.



Layer 2 Spanned EtherChannels are not supported.

Policy-based routing

When using individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Policy-Based Routing (PBR).

We recommend this method if you are already using PBR, and want to take advantage of your existing infrastructure.

PBR makes routing decisions based on a route map and ACL. You must manually divide traffic between all Firewall Threat Defenses in a cluster. Because PBR is static, it may not achieve the optimum load balancing result at all times. To achieve the best performance, we recommend that you configure the PBR policy so that forward and return packets of a connection are directed to the same Firewall Threat Defense.

For example, if you have a Cisco router, redundancy can be achieved by using Cisco IOS PBR with Object Tracking. Cisco IOS Object Tracking monitors each Firewall Threat Defense using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular Firewall Threat Defense.

Equal-cost multi-path routing

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Equal-Cost Multi-Path (ECMP) routing.

We recommend this method if you are already using ECMP, and want to take advantage of your existing infrastructure.

ECMP routing considerations

ECMP routing can forward packets over multiple best paths that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses and/or source and destination ports can be used to send a packet to one of the next hops.

If you use static routes for ECMP routing, then the Firewall Threat Defense failure can cause problems; the route continues to be used, and traffic to the failed Firewall Threat Defense will be lost.

If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each Firewall Threat Defense to participate in dynamic routing.

Cluster Control Link

Each node must dedicate one interface as a VXLAN (VTEP) interface for the cluster control link. For more information about VXLAN, see Configure VXLAN interfaces, on page 1071 .

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

VTEP Source Interface

The VTEP source interface is a regular Firewall Threat Defense Virtual interface with which you plan to associate the VNI interface. You can configure one VTEP source interface to act as the cluster control link. The source interface is reserved for cluster control link use only. Each VTEP source interface has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the cluster control link interfaces.

VNI Interface

A VNI interface is similar to a VLAN interface. It is a virtual interface that keeps network traffic separated on a given physical interface by using tagging. You can only configure one VNI interface. Each VNI interface has an IP address on the same subnet.

Peer VTEPs

Unlike regular VXLAN for data interfaces, which allows a single VTEP peer, Firewall Threat Defense Virtual clustering allows you to configure multiple peers.

Cluster control link traffic overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

Configuration replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Management network

You must manage each node using the Management interface. Management from a data interface is not supported with clustering.

Licenses for clustering: Private cloud

Performance tier license requirements

Each Firewall Threat Defense Virtual cluster node requires the same performance tier license. We recommend using the same number of CPUs and memory for all members, or else performance will be limited on all nodes to match the least capable member. The throughput level will be replicated from the control node to each data node so they match.

Feature license assignment

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the Cloud-Delivered Firewall Management Center, you can specify the feature licenses you want to use for the cluster. Before you create the cluster, it doesn't matter which licenses are assigned to the data nodes; the license settings for the control node are replicated to each of the data nodes. You can modify licenses for the cluster by clicking **Edit Licenses** in **System** (⚙) > **Licenses** > **Smart Licenses** or choosing **Devices** > **Device Management**, clicking **Edit** (✎) for the cluster, and then in the **License** area, clicking **Edit** (✎).

Note If you add the cluster before the Cloud-Delivered Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Cloud-Delivered Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Note APEX, PLUS, and VPN-Only licenses are not supported for clustering. Before you create a cluster, remove unsupported license assignments from the devices that will become cluster nodes.

Prerequisites for clustering: Private cloud

This reference provides the complete set of requirements that must be met before configuring clustering for threat defense virtual devices in private cloud deployments.

Model requirements

- FTDv5, FTDv10, FTDv20, FTDv30, FTDv50, FTDv100
- VMware or KVM
- A maximum of 4 nodes in a cluster in a 2x2 configuration is supported. You can set up a maximum of two hosts with a maximum of two threat defense virtual instances in each host.

Hardware and software requirements

All units in a cluster:

- Must have jumbo frame reservation enabled for the cluster control link. Do this in the Day 0 configuration when you deploy the Firewall Threat Defense Virtual by setting `"DeploymentType": "Cluster"`. Otherwise, you must restart each node to enable jumbo frames after the cluster has formed and is healthy.
- (KVM only) Must use CPU hard partitioning (CPU pinning) for all VMs on the KVM host.
- Must be the same performance tier. We recommend using the same number of CPUs and memory for all nodes, or performance will be limited on all nodes to match the least capable node.
- Must use the management interface for Cloud-Delivered Firewall Management Center communications. Data interface management is not supported.
- Must run the same version, except during upgrade. Hitless upgrade is supported.
- Must not have any deployment pending or in progress.
- Must not have any unsupported features configured on the control node: Unsupported features and clustering, on page 832.
- Must not have VPN configured on the data nodes. The control node can have site-to-site VPN configured.

Note Devices must not be in transparent mode for cluster configuration.

Cloud-Delivered Firewall Management Center Requirements

Make sure the Cloud-Delivered Firewall Management Center NTP server is set to a reliable server that is reachable by all cluster nodes to ensure proper clock sync. By default, the device uses the same NTP server

as the Cloud-Delivered Firewall Management Center. If the time is not set to be the same on all cluster nodes, they can be removed automatically from the cluster.

Switch requirements

Be sure to complete the switch configuration before you configure clustering. Make sure the ports connected to the cluster control link have the correct (higher) MTU configured. By default, the cluster control link MTU is set to 154 bytes higher than the data interfaces. If the switches have an MTU mismatch, the cluster formation will fail. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the initial ping fails, the node tries a ping using a smaller packet size (the MTU divided by 2, then by 4, then by 8) until a ping succeeds. A notification is generated so you can fix the MTU mismatch on connecting switches and try again.

Guidelines for clustering: Private cloud

High availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Additional clustering guidelines

Consider these additional guidelines when configuring clustering:

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Firewall Threat Defense Virtual, adding an additional switch to form a VSS or vPC) you should disable the health check feature and also disable interface monitoring for the interfaces that are affected by the topology changes. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.
- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.
- We do not support VXLAN for data interfaces; only the cluster control link supports VXLAN.

Default cluster configurations

Consider these default configurations when setting up clustering:

- The cLACP system ID is auto-generated, and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.

- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Upgrade guidelines

Configure clustering: private cloud

To configure clustering after you deploy your Firewall Threat Defense Virtuals, perform the following tasks.

Add nodes to the Management Center

This task allows you to prepare cluster nodes for clustering configuration by adding them to the management center as individual devices.

Before configuring clustering, deploy each cluster node, then add the devices as standalone units on the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Deploy each cluster node by following the procedures given in the Secure Firewall Threat Defense Virtual getting started guides.

All units in a cluster:

- Must have jumbo frame reservation enabled for the cluster control link. Do this in the Day 0 configuration when you deploy the Firewall Threat Defense Virtual by setting `"DeploymentType": "Cluster"`. Otherwise, you must restart each node to enable jumbo frames after the cluster has formed and is healthy.
- (KVM only) Must use CPU hard partitioning (CPU pinning) for all VMs on the KVM host.

Step 2 Add each node to the Cloud-Delivered Firewall Management Center as a standalone device in the same domain and group.

You can create a cluster with a single device, and then add more nodes later. The initial settings (licensing, access control policy) that you set when you add a device will be inherited by all cluster nodes from the control node. You will choose the control node when forming the cluster.

Create a cluster

Creating a cluster allows multiple devices to work together as a single logical unit, providing high availability and increased throughput capacity.

Form a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Before you begin

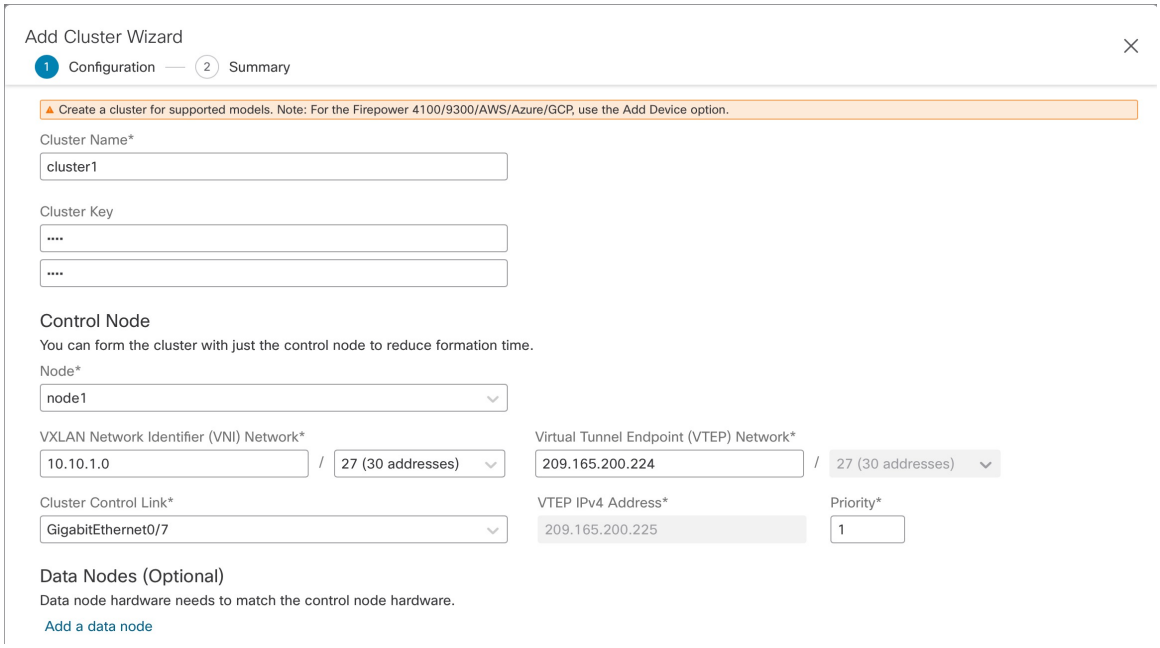
Some features are not compatible with clustering. Wait to configure until clustering is enabled. Some features will block cluster creation if they are already configured. For example, do not configure any IP addresses on interfaces, or unsupported interface types such as BVIs.

Procedure

Step 1 Choose **Devices** > **Device Management**, and then choose **Add** > **Cluster**.

The **Add Cluster Wizard** appears.

Figure 141: Add Cluster Wizard



Step 2 Specify a **Cluster Name** and an authentication **Cluster Key** for control traffic.

- **Cluster Name**—An ASCII string of 1 to 38 characters. The cluster name can contain only letters, digits, underscores (_), and hyphens (-).
- **Cluster Key**—An ASCII string of 1 to 63 characters. The cluster key cannot be a single digit and must not contain spaces. The **Cluster Key** value is used to generate the encryption key. This encryption does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in clear text.

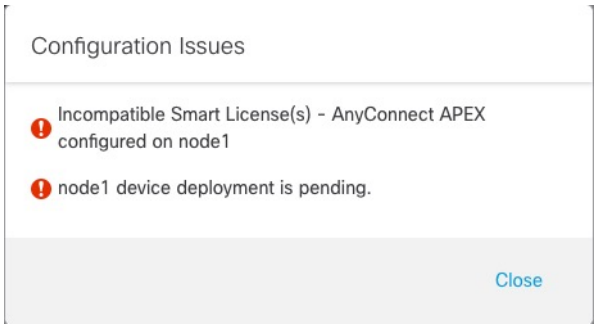
Step 3 For the **Control Node**, set the following:

- **Node**—Choose the device that you want to be the control node initially. When the Cloud-Delivered Firewall Management Center forms the cluster, it will add this node to the cluster first so it will be the control node.

  Note

If you see an **Error** ( ) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation. For example:

Figure 142: Configuration Issues



To resolve the above issues, remove the unsupported VPN license and deploy pending configuration changes to the device.

- **VXLAN Network Identifier (VNI) Network**—Specify an IPv4 subnet for the VNI network; IPv6 is not supported for this network. Specify a **24**, **25**, **26**, or **27** subnet. An IP address will be auto-assigned to each node on this network. The VNI network is the encrypted virtual network that runs on top of the physical VTEP network.
- **Cluster Control Link**—Choose the physical interface you want to use for the cluster control link.
- **Virtual Tunnel Endpoint (VTEP) Network**—Specify an IPv4 subnet for the physical interface network; IPv6 is not supported for this network. The VTEP network is a different network than the VNI network, and it is used for the physical cluster control link.
- **VTEP IPv4 Address**—This field will be auto-populated with the first address on the VTEP network.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority. Even if you set the priority to be lower than other nodes, this node will still be the control node when the cluster is first formed.

Step 4 For **Data Nodes (Optional)**, click **Add a data node** to add a node to the cluster.

You can form the cluster with only the control node for faster cluster formation, or you can add all nodes now. Set the following for each data node:

- **Node**—Choose the device that you want to add.

  Note

  If you see an **Error** ( ) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation.

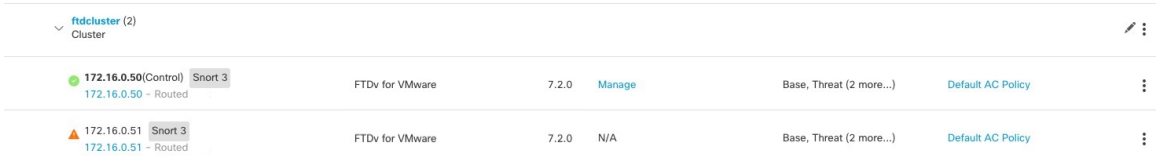
- **VTEP IPv4 Address**—This field will be auto-populated with the next address on the VTEP network.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority.

Step 5 Click **Continue**. Review the **Summary**, and then click **Save**.

The cluster bootstrap configuration is saved to the cluster nodes. The bootstrap configuration includes the VXLAN interface used for the cluster control link.

The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster nodes.

Figure 143: Cluster Management

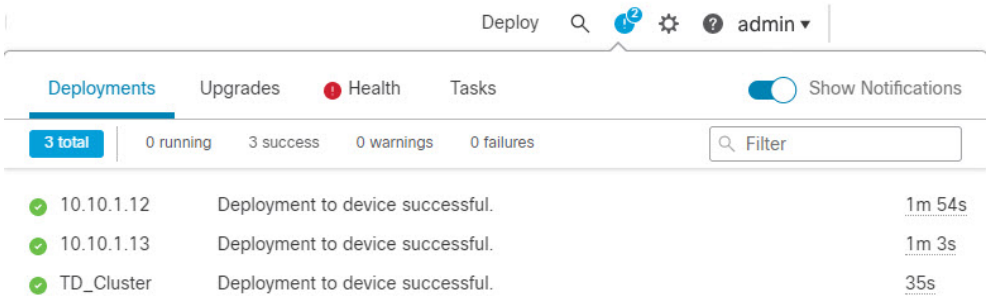


A node that is currently registering shows the loading icon.

Figure 144: Node Registration



You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**. The Cloud-Delivered Firewall Management Center updates the Cluster Registration task as each node registers.

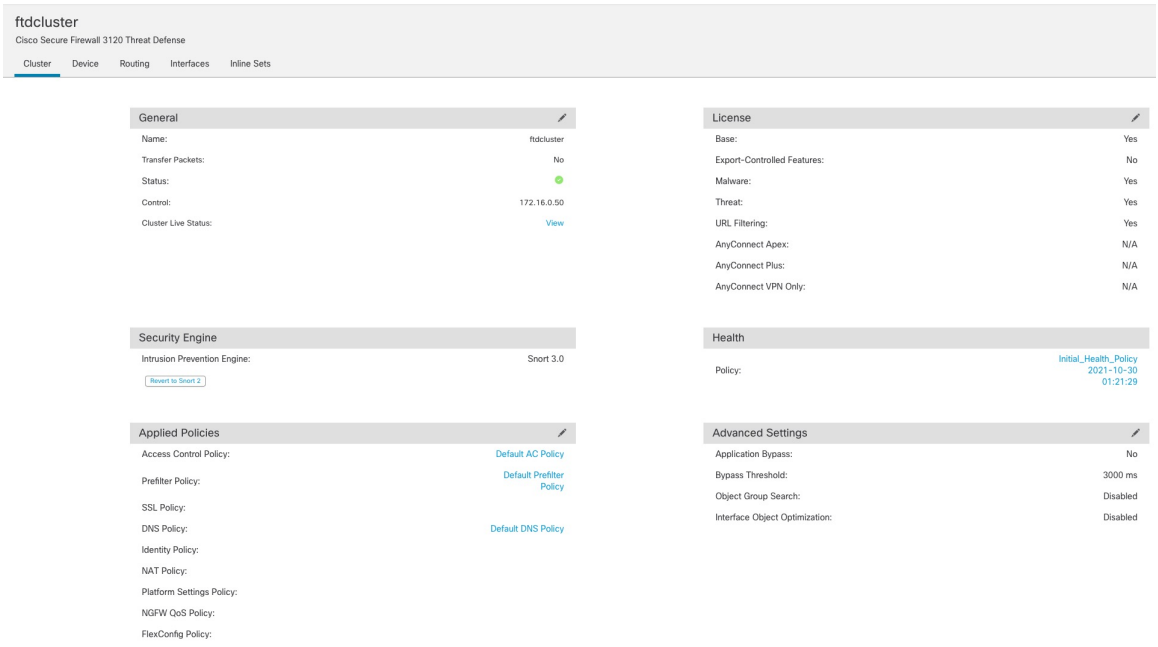


Step 6 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

Most configuration can be applied to the cluster as a whole, and not nodes in the cluster. For example, you can change the display name per node, but you can only configure interfaces for the whole cluster.

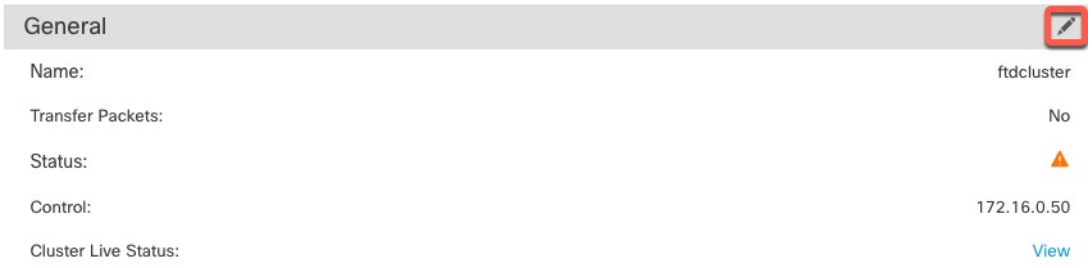
Step 7 On the **Devices** > **Device Management**, **Cluster** screen, you see **General** and other settings for the cluster.

Figure 145: Cluster Settings

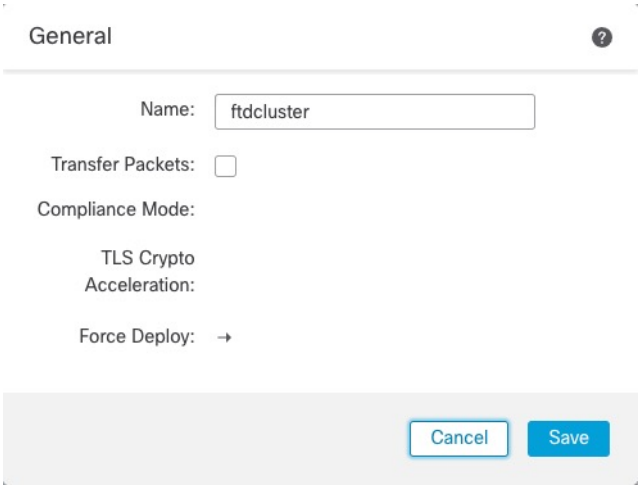


See the following cluster-specific items in the **General** area:

- **General > Name**—Change the cluster display name by clicking the **Edit** (✎).



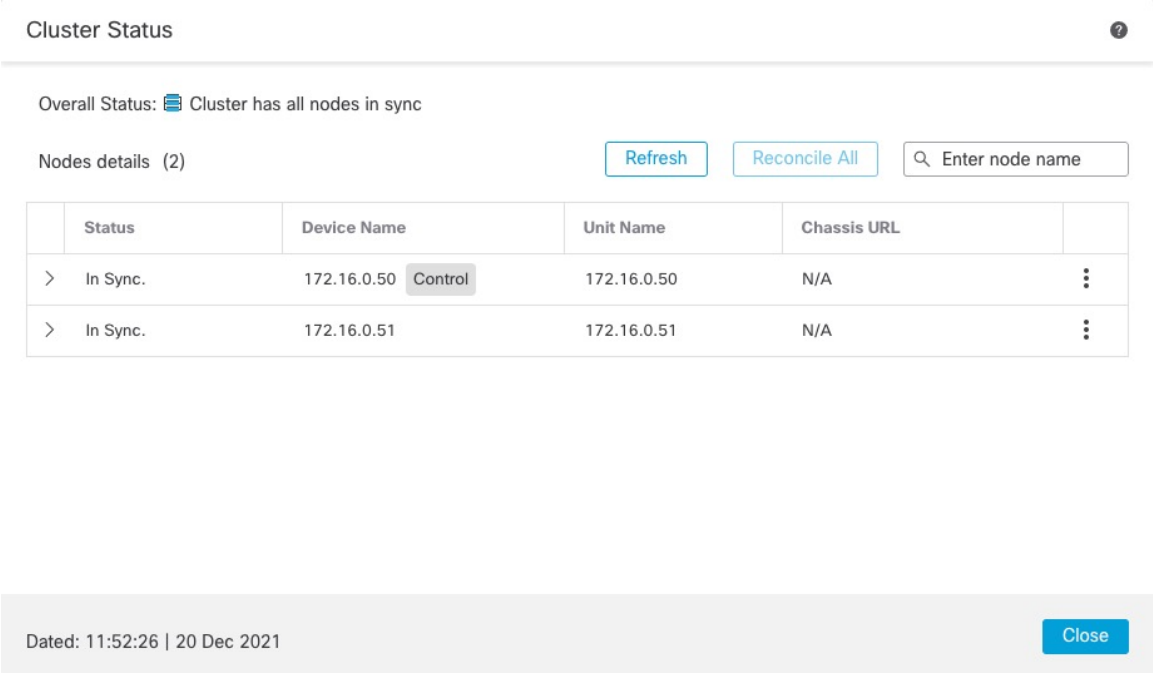
Then set the **Name** field.



- **General > View**—Click the **View** link to open the **Cluster Status** dialog box.



The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile All**.



Step 8 On the **Devices** > **Device Management** and then click **Add** > **Device**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.

Figure 146: Device Settings

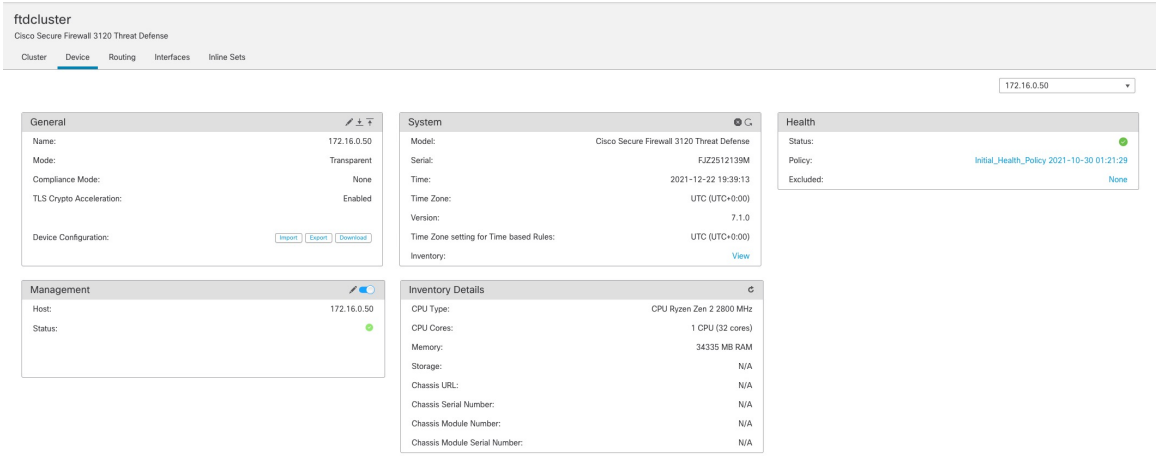
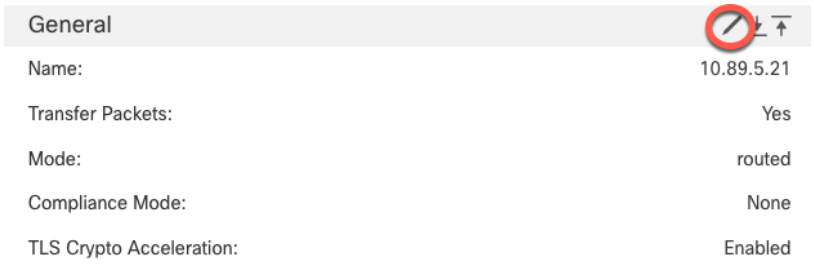


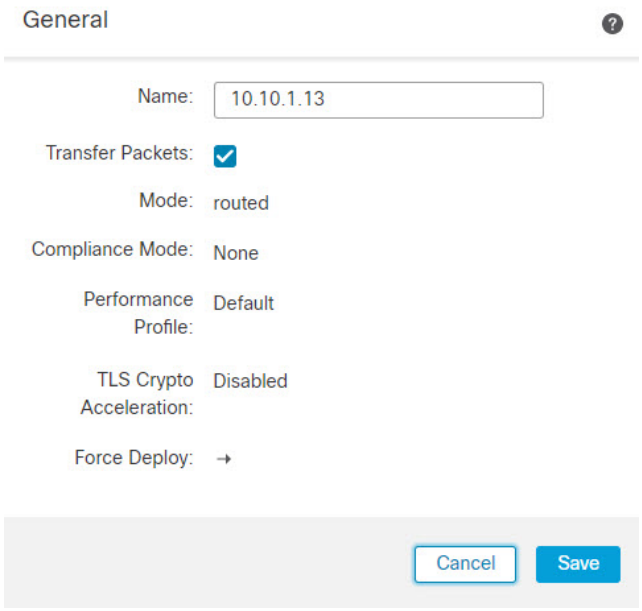
Figure 147: Choose Node



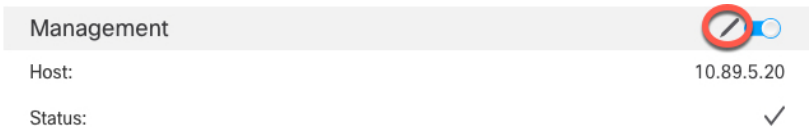
- **General** > **Name**—Change the cluster member display name by clicking the **Edit** ( ).



Then set the **Name** field.



- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the Cloud-Delivered Firewall Management Center so that it can reach the device on the network. First disable the connection, edit the **Host** address in the **Management** area, then re-enable the connection.



Step 9 If you deployed your cluster nodes without enabling jumbo-frame reservation, then restart all cluster nodes to enable jumbo frames, which are required for the cluster control link. See Shut down or restart the device, on page 107.

If you previously enabled jumbo-frame reservation, you can skip this step.

Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead (100 bytes) and VXLAN overhead (54 bytes). When you create the cluster, the MTU is set to 154 bytes higher than the highest data interface MTU (1654 by default). If you later increase the data interface MTU, be sure to also increase the cluster control link MTU. For example, because the maximum MTU is 9198 bytes, then the highest data interface MTU can be 9044, while the cluster control link can be set to 9198. See Configure the MTU, on page 1109.

Make sure you configure switches connected to the cluster control link to the correct (higher) MTU; otherwise, cluster formation might fail.

Configure interfaces

Configure interfaces to be compatible with clustering by setting up individual interfaces that are normal routed interfaces, each with their own IP address taken from a pool of IP addresses.

Individual interfaces are normal routed interfaces, each with their own IP address taken from a pool of IP addresses. The Main cluster IP address is a fixed address for the cluster that always belongs to the current control node. All data interfaces must be Individual interfaces.

Note You cannot use subinterfaces.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Address Pools** to add an IPv4 address pool, an IPv6 address pool, or both. See Configure address pools, on page 2954.

Include at least as many addresses as there are units in the cluster. The Virtual IP address is not a part of this pool, but needs to be on the same network. You cannot determine the exact Local address assigned to each unit in advance.

Step 2 Choose **Devices** > **Device Management**, and click **Edit** ( ) next to the cluster.

Step 3 Click **Interfaces**, and then click **Edit** ( ) for a data interface.

Step 4 On the **IPv4**, enter the **IP Address** and mask. This IP address is a fixed address for the cluster, and always belongs to the current control unit.

Step 5 From the **IPv4 Address Pool** drop-down list, choose the address pool you created.

Note
If you want to manually assign a MAC address to this interface, you need to create a **MAC-address pool** using FlexConfig.

Step 6 On **IPv6** > **Basic**, from the **IPv6 Address Pool** drop-down list, choose the address pool you created.

Step 7 Configure other interface settings as normal.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Cluster Health Monitor Settings

This task allows you to modify cluster health monitor settings to control how cluster nodes monitor system health and automatically rejoin the cluster after failures.

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings for monitoring cluster node health. You can monitor any port-channel ID, any single physical interface ID, and the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link. It is always monitored.

When a topology change occurs (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature. Also, disable interface monitoring for any interfaces that are disabled. When

the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Procedure

Step 1 Choose **Devices** > **Device Management**.

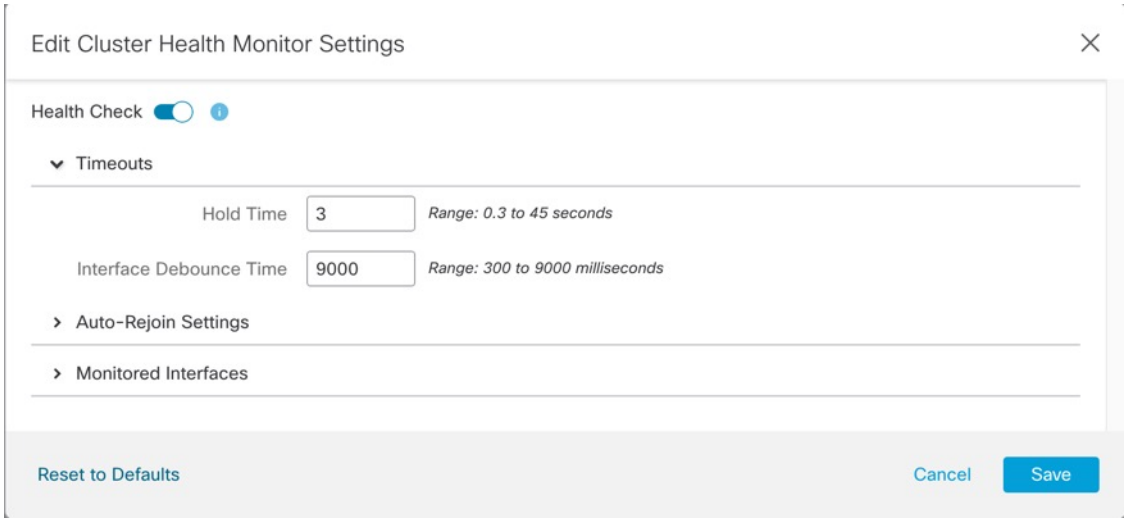
Step 2 Next to the cluster you want to modify, click **Edit** ( ).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** ( ).

Step 5 Disable the system health check by clicking the **Health Check** slider.

Figure 148: Disable the system health check



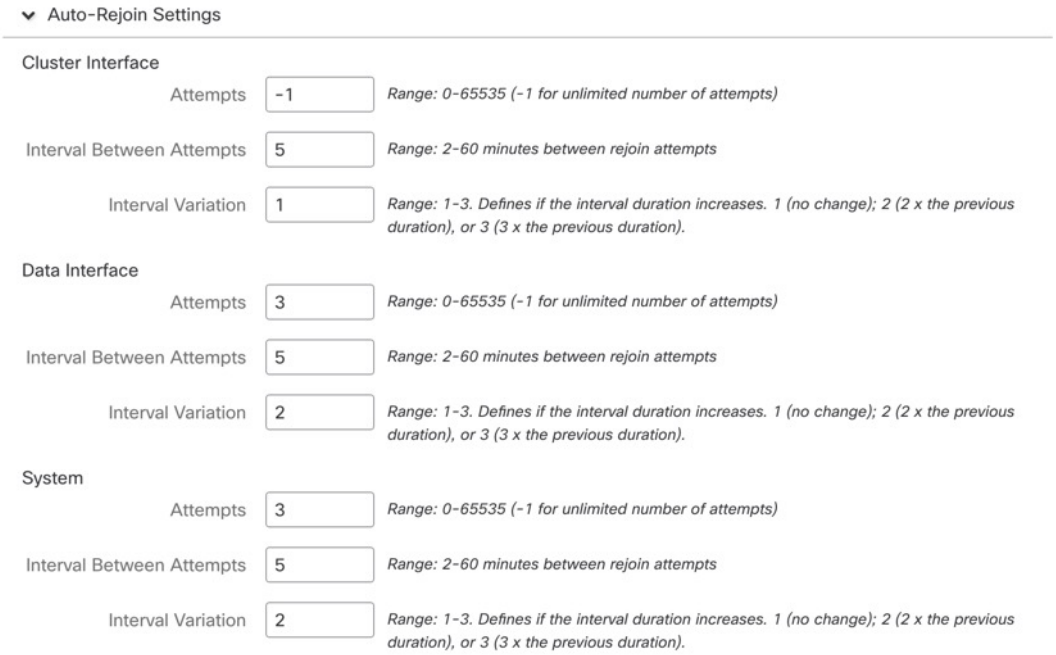
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between 0.3 and 45 seconds. The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the specified number of milliseconds before marking the interface as failed and removing the node from the cluster. If an EtherChannel transitions from a down state to an up state, such as after the switch reloads or enables an EtherChannel, a longer debounce time can prevent the interface from appearing failed on a cluster node if another cluster node bundles the ports more quickly.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 149: Configure Auto-Rejoin settings

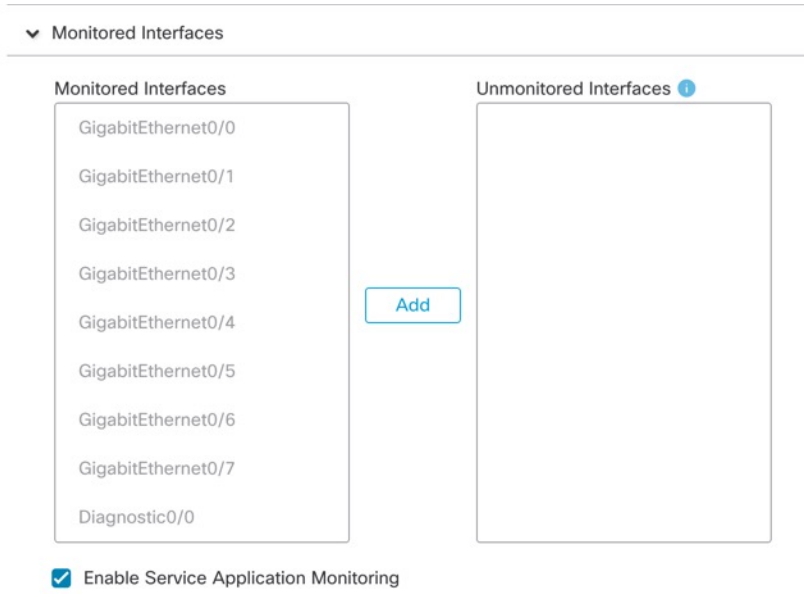


Set values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 150: Configure monitored interfaces



The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The time required to remove a member from the cluster depends on the type of interface and whether the node is established or is joining the cluster. By default, health check is enabled for all interfaces and the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Deploy configuration changes.

Manage cluster nodes

Describes how to manage cluster nodes to maintain high availability and optimal performance across your Cisco Secure Firewall cluster deployment.

Add a new cluster node

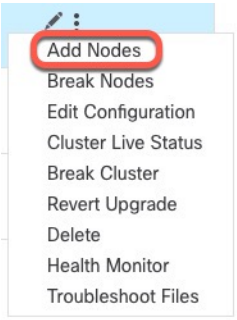
Adding new cluster nodes expands cluster capacity and provides additional redundancy for your security infrastructure.

You can add one or more new cluster nodes to an existing cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) for the cluster, and choose **Add Nodes**.

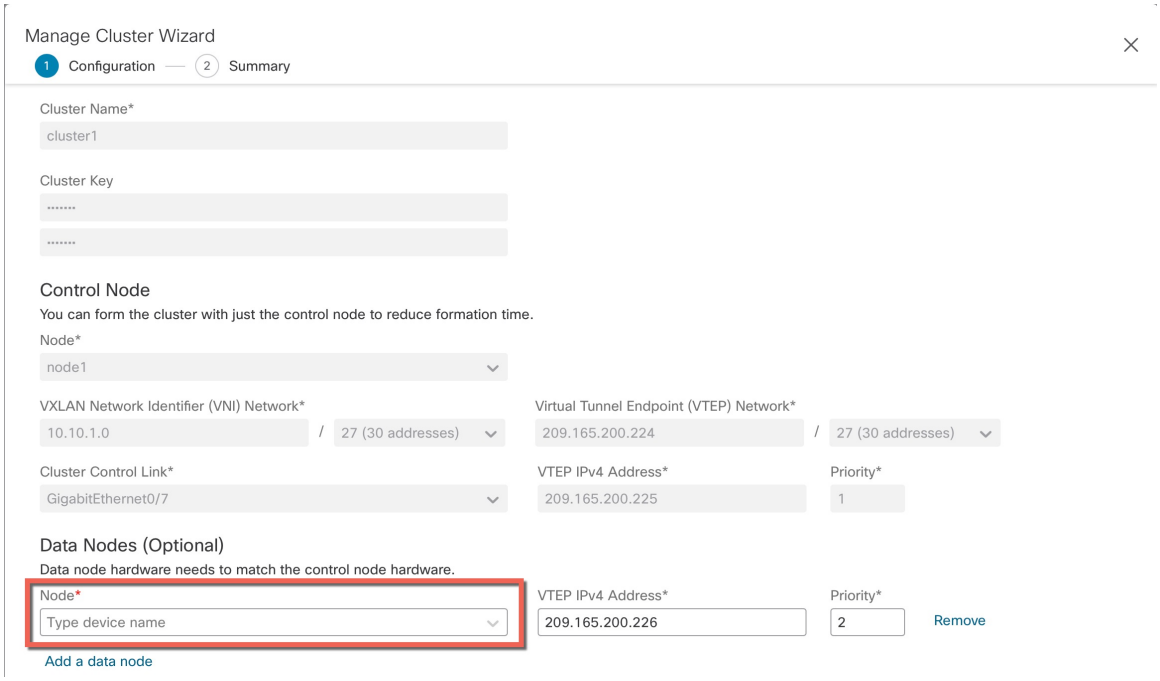
Figure 151: Add nodes



The **Manage Cluster Wizard** appears.

Step 2 From the **Node** menu, choose a device, and adjust the IP address and priority if desired.

Figure 152: Manage cluster wizard

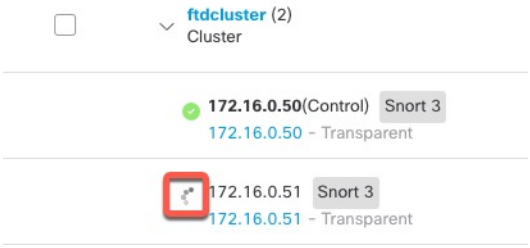


Step 3 To add additional nodes, click **Add a data node**.

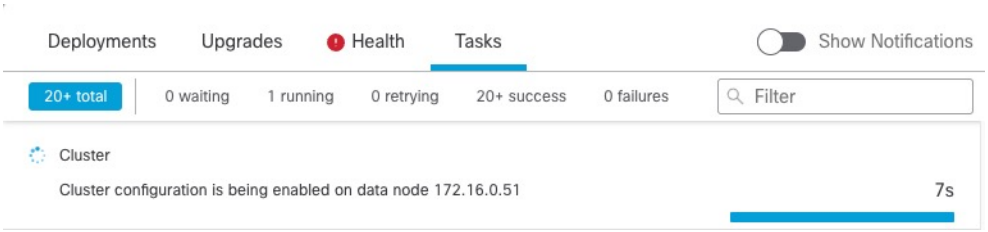
Step 4 Click **Continue**. Review the **Summary**, and then click **Save**

The node that is currently registering shows the loading icon.

Figure 153: Node registration



You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**.



Break a node

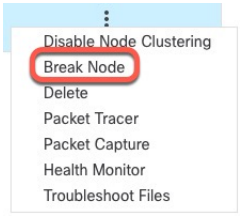
Breaking a node removes it from the cluster configuration, converting it to a standalone device.

You can remove a node from the cluster so that it becomes a standalone device. You cannot break the control node unless you break the entire cluster. The data node has its configuration erased.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) for the node you want to break, and choose **Break Node**.

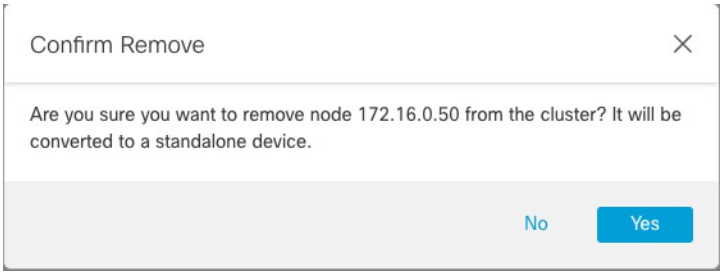
Figure 154: Break a node



You can optionally break one or more nodes from the cluster More menu by choosing **Break Nodes**.

Step 2 You are prompted to confirm the break. Click **Yes**.

Figure 155: Confirm break



You can monitor the cluster node break by clicking the **Notifications** icon and choosing **Tasks**.

Break the cluster

Break the cluster to convert all nodes to standalone devices when clustering is no longer needed.

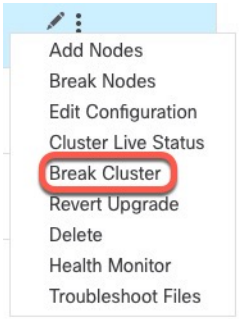
You can break the cluster and convert all nodes to standalone devices. The control node retains the interface and security policy configuration, while data nodes have their configuration erased.

Procedure

Step 1 Make sure all cluster nodes are being managed by the Cloud-Delivered Firewall Management Center by reconciling nodes. Refer to Reconcile cluster nodes, on page 824.

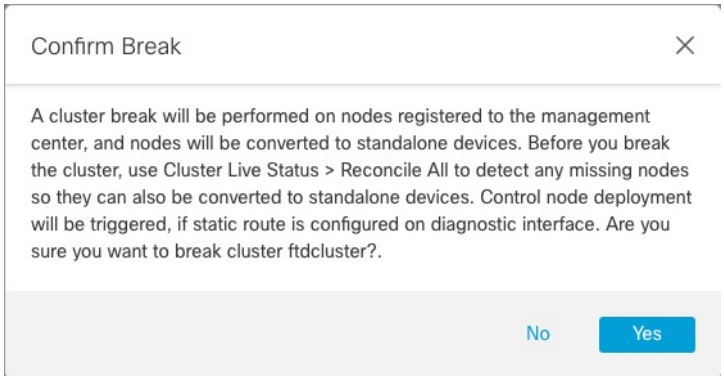
Step 2 Choose **Devices** > **Device Management**, click the **More** (⋮) for the cluster, and choose **Break Cluster**.

Figure 156: Break cluster



Step 3 You are prompted to break the cluster. Click **Yes**.

Figure 157: Confirm break



You can monitor the cluster break by clicking the **Notifications** icon and choosing **Tasks**.

Disable clustering

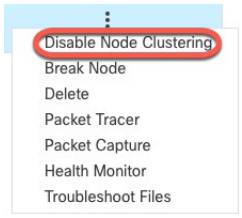
Temporarily deactivate a node in a cluster while keeping it visible in the device list.

You may want to deactivate a node in preparation for deleting the node, or temporarily for maintenance. This procedure is meant to temporarily deactivate a node; the node will still appear in the Cloud-Delivered Firewall Management Center device list. When a node becomes inactive, all data interfaces are shut down.

Procedure

Step 1 For the unit you want to disable, choose **Devices** > **Device Management**, click **More** (⋮), and choose **Disable Node Clustering**.

Figure 158: Disable clustering



If you disable clustering on the control node, one of the data nodes will become the new control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to reestablish the connections on the new control node. You cannot disable clustering on the control node if it is the only node in the cluster.

Step 2 Confirm that you want to disable clustering on the node.

The node will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.

Step 3 To reenable clustering, refer to Rejoin the cluster, on page 822.

Rejoin the cluster

Rejoin a node to the cluster after it was removed because of failure or when clustering was manually disabled.

If a node was removed from the cluster, for example, due to a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management**, click **More** (⋮), and choose **Enable Node Clustering**.

Step 2 Confirm that you want to enable clustering on the node.

Change the control node

This task changes the control node in a cluster when you need to specify the exact unit you want to become the control node.

The best method to change the control node is to disable clustering on the control node, wait for a new control election, and then re-enable clustering. Use this procedure when you must specify the exact unit you want to become the control node.



Caution If you must specify the *exact* unit you want to become the control node, use the procedure in this section. Note that for centralized features, if you force a control node change using either method, then all connections are dropped, and you have to re-establish the connections on the new control node.

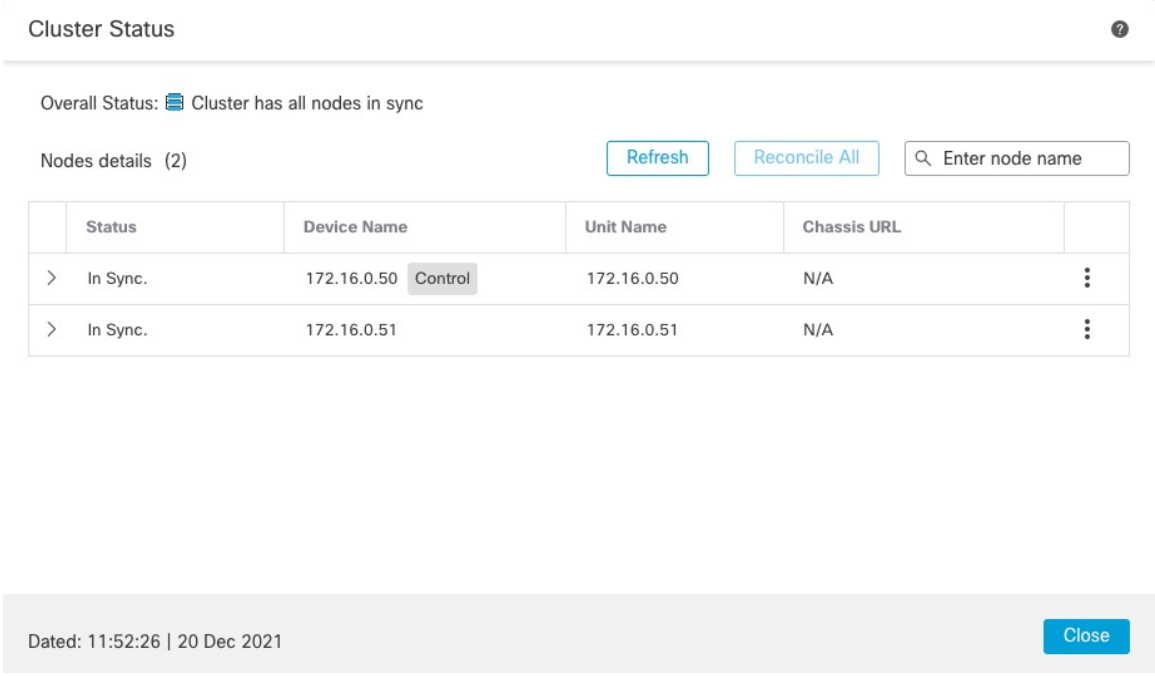
Before you begin

Follow these steps to change the control node:

Procedure

Step 1 Open the **Cluster Status** dialog box by choosing **Devices** > **Device Management**, **More** (⋮) > **Cluster Live Status**.

Figure 159: Cluster status



Step 2 For the unit you want to become the control unit, choose **More** (⋮) > **Change Role to Control**.

Step 3 You are prompted to confirm the role change. Select the checkbox and click **OK**.

Edit the cluster configuration

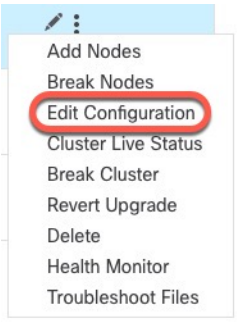
Edit the cluster configuration to modify cluster settings and parameters as needed for your network environment.

You can edit the cluster configuration. If you change any values other than the VTEP IP address for a node or node priority, the cluster will be broken and reformed automatically. Until the cluster is reformed, you may experience traffic disruption. If you change the VTEP IP address for a node or node priority, only the affected nodes are broken and readded to the cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) for the cluster, and choose **Edit Configuration**.

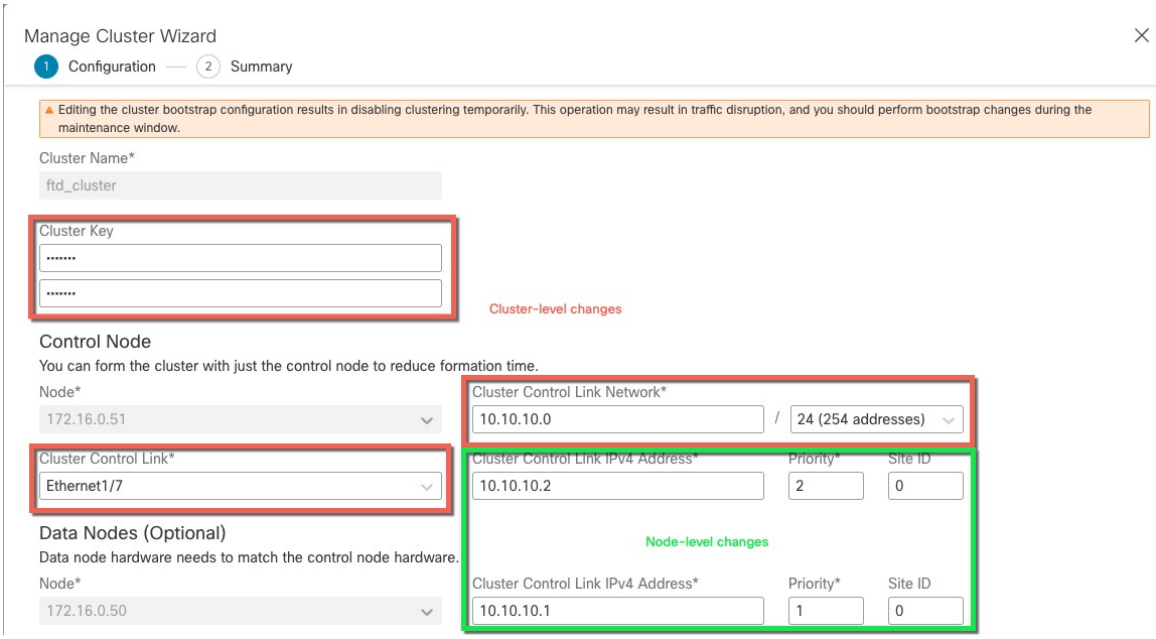
Figure 160: Edit configuration



The **Manage Cluster Wizard** appears.

Step 2 Update the cluster configuration.

Figure 161: Manage cluster wizard



Step 3 Click **Continue**. Review the **Summary**, and then click **Save**.

Reconcile cluster nodes

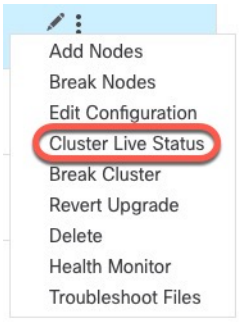
Use this procedure to reconcile cluster membership when a cluster node fails to register with the Cloud-Delivered Firewall Management Center.

If a cluster node fails to register, you can reconcile the cluster membership from the device to the Cloud-Delivered Firewall Management Center. For example, a data node might fail to register if the Cloud-Delivered Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

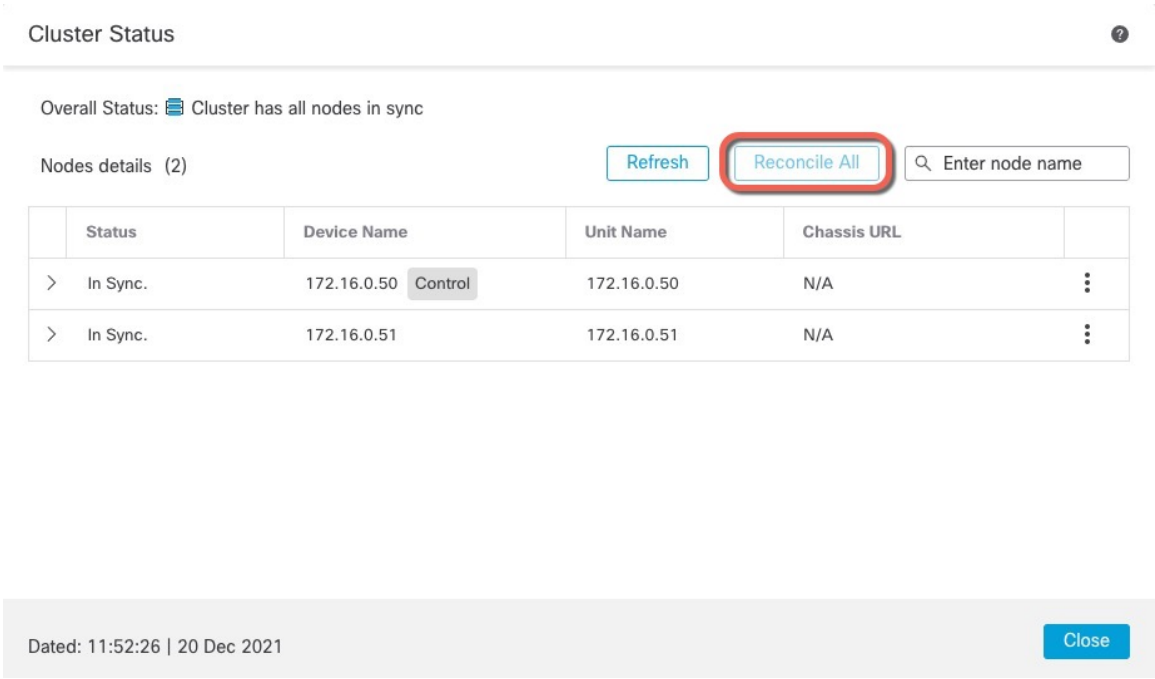
Step 1 Choose **Devices** > **Device Management**, **More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Figure 162: Cluster live status



Step 2 Click **Reconcile All**.

Figure 163: Reconcile all



Delete the cluster or nodes from the Management Center

Delete clusters or nodes from Security Cloud Control to manage cluster membership and device visibility.

You can delete the cluster from the Security Cloud Control, which keeps the cluster intact. You might want to delete the cluster if you want to add the cluster to a new Security Cloud Control.

You can also delete a node from the Security Cloud Control without breaking the node from the cluster. Although the node is not visible in the Security Cloud Control, it is still part of the cluster, and it will continue to pass traffic and could even become the control node. You cannot delete the current control node. You might want to delete the node if it is no longer reachable from the Security Cloud Control, but you still want to keep it as part of the cluster.

Procedure

Step 1 Log into Security Cloud Control and click **Security Devices**.

Step 2 Click the **FTD** tab and locate the cluster you want. Select it so the device row is highlighted.

Step 3 Do the following:

- To delete a node within the cluster, in the **Cluster** pane on the right, click the delete icon next to the device you want to delete.
- To delete the cluster, in the **Device Actions** pane on the right, click **Remove**.

Step 4 When prompted, select **OK** to confirm the removal of the selected device.

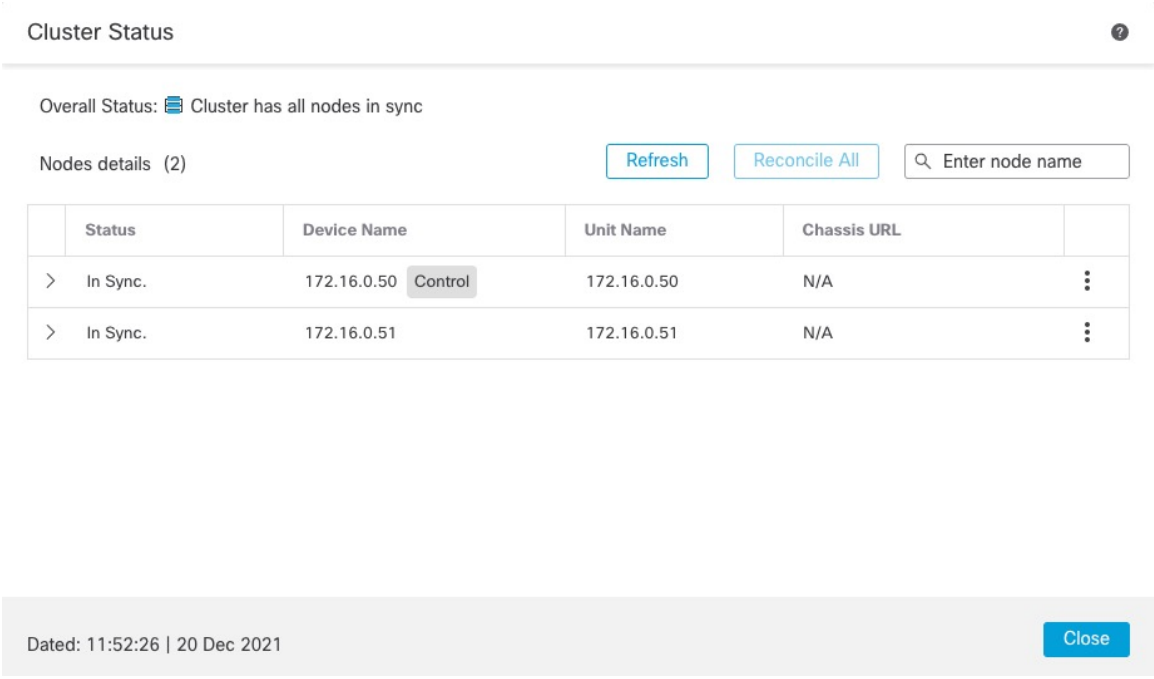
Monitor cluster status

You can monitor the cluster in the Cloud-Delivered Firewall Management Center and at the Firewall Threat Defense CLI.

Monitor cluster status through these methods:

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management**, **More** (⋮) icon or from the **Devices** > **Device Management**, click **Add**, choose the **Cluster** page **General** area **Cluster Live Status** link.

Figure 164: Cluster status



The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes these states:

- In Sync—The node is registered with the Cloud-Delivered Firewall Management Center.
- Pending Registration—The node is part of the cluster, but has not yet registered with the Cloud-Delivered Firewall Management Center. If a node fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The node is registered with the Cloud-Delivered Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the node from the cluster.
- Joining cluster...—The node is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Cloud-Delivered Firewall Management Center.

For each node, you can view the **Summary** or the **History**.

Figure 165: Node summary

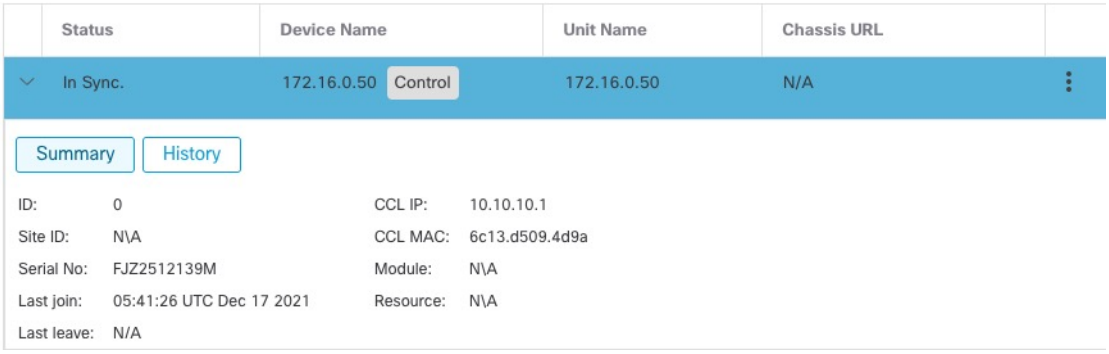
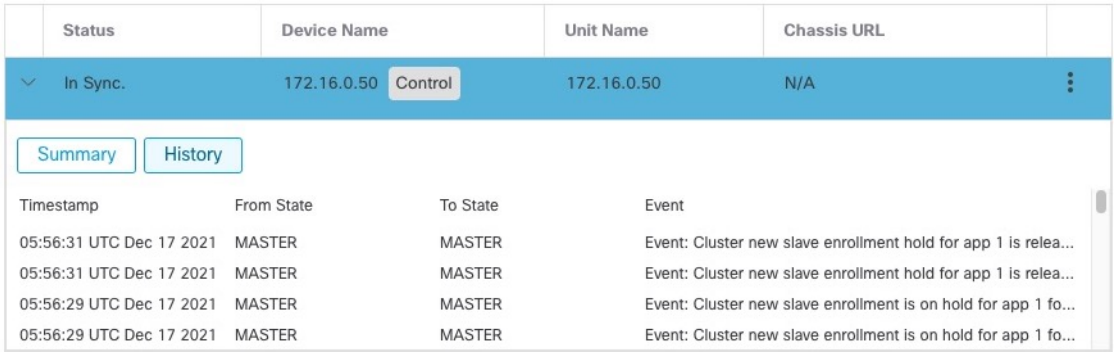


Figure 166: Node history



- **System** (⚙) > **Tasks** page.

  The **Tasks** page shows updates of the Cluster Registration task as each node registers.

- **Devices** > **Device Management** and then click **Add** > **Device** *cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member nodes, including the control node shown with its role next to the IP address. For nodes that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Cluster health monitor

When a Firewall Threat Defense is the control node of a cluster, the Cloud-Delivered Firewall Management Center collects various metrics periodically from the device metric data collector.

The cluster health monitor is comprised of these components:

- Overview dashboard: Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Cloud-Delivered Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.



Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard: Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.

- Member Performance dashboard: Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard: Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links: Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range: An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard: Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

View the cluster health monitor

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

Ensure you have created a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Follow these steps to view the cluster health monitor:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed cluster devices.

Step 3 Click on the cluster name to view the cluster health statistics.

The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution: Traffic and packet distribution across the cluster nodes.
- Member Performance: Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL: Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 View device monitor health and performance metrics in several predefined dashboards.

The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.

- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics that are related to the Snort process.
- ASP drops: Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. Refer to Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 54: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number
Packets	Packet distribution count in the cluster for every second.	number

Reference for clustering

This section includes more information about how clustering operates.

Threat Defense features and clustering

Some Firewall Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported features and clustering

Centralized features for clustering

These features are only supported on the control node, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member nodes to the control node over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control nodes before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control node.

For centralized features, if the control node fails, all connections are dropped, and you have to re-establish the connections on the new control node.

Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, refer to the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Cloud-Delivered Firewall Management Center GUI. For more information, refer to FlexConfig policies, on page 1309.

- These application inspections are supported:
  - DCERPC and ESMTP
  - NetBIOS and PPTP
  - RSH
  - SQLNET and SUNRPC
  - TFTP
  - XDMCP
- Static route monitoring
- Site-to-site VPN

- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Dynamic routing

Connection settings and clustering

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

Dynamic routing and clustering

Dynamic routing and clustering are network configuration approaches that

- enable each node to run the routing protocol as a standalone router in Individual interface mode,
- allow routes to be learned by each node independently, and
- require separate router IDs for each node when communicating with external routers.

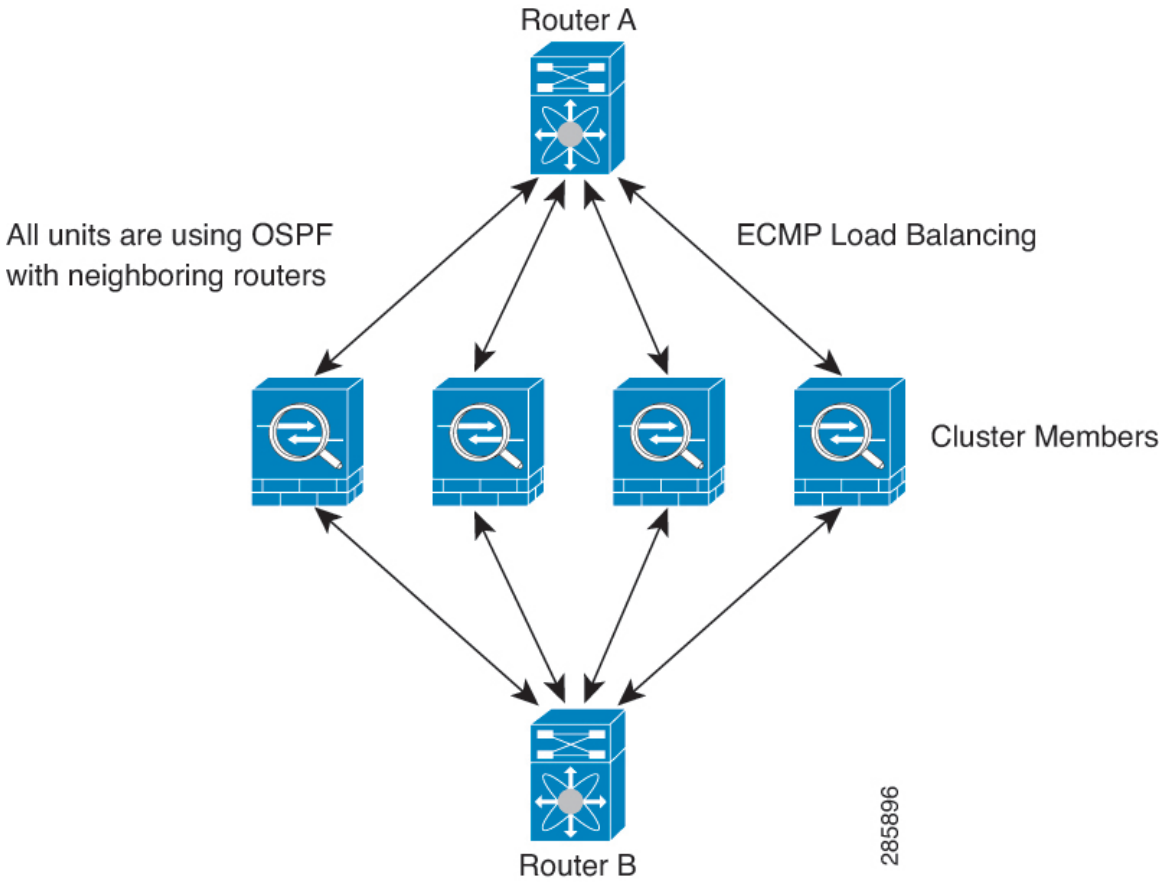
Routing behavior and load balancing

In Individual interface mode, ECMP is used to load balance traffic between multiple equal-cost paths. External routers can discover multiple paths to destinations through different cluster nodes.

Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when communicating with external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

Figure 167: Dynamic routing in individual interface mode



FTP and clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

NAT and clustering

Provides information about NAT behavior and limitations in clustering deployments.

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different Firewall Threat Defenses in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the Firewall Threat Defense that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- No Proxy ARP—For individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with an ASA that may no longer be in the cluster. The upstream router requires a static route or policy-based routing (PBR) with Object Tracking for the mapped addresses, which must point to the main cluster IP address.
- No interface PAT on an individual interface—Interface PAT is not supported for Individual interfaces.
- PAT with Port Block Allocation—See the following guidelines for this feature:
  - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
  - Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
  - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
  - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.
- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
  - FTP

- RSH
- SQLNET
- TFTP
- XDMCP
- SIP

- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

SIP inspection and clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and clustering

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

Cisco trustsec and clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

VPN and clustering

VPN functionality is limited to the control node and does not take advantage of the cluster high availability capabilities. If the control node fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control node is elected, you must reestablish the VPN connections.

For connections to an individual interface when using Policy-Based Routing (PBR) or Equal-Cost Multi-Path (ECMP), you must always connect to the Main cluster IP address, not a Local address.

VPN-related keys and certificates are replicated to all nodes.

Note Remote access VPN is not supported with clustering.

Performance scaling factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

Throughput calculation

For example, if your model can handle approximately 10 Gbps of traffic when running alone, then for a cluster of 8 units, the maximum combined throughput will be approximately 80% of 80 Gbps (8 units x 10 Gbps): 64 Gbps.

Control node election

The key components involved in the control node election process are:

- Cluster nodes: Communicate over the cluster control link and participate in the election process
- Election requests: Broadcast every three seconds by nodes when clustering is enabled
- Priority values: Set between 1 and 100 where 1 is the highest priority
- Control node: The elected node that coordinates cluster operations

These stages describe how cluster nodes elect a control node:

1. When clustering is enabled for a node (or when it first starts up with clustering already enabled), the node broadcasts an election request every three seconds.
2. Any other nodes with a higher priority respond to the election request. The priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a node does not receive a response from another node with a higher priority, then it becomes the control node.

Note If multiple nodes tie for the highest priority, the cluster node name and then the serial number is used to determine the control node.

4. If a node later joins the cluster with a higher priority, it does not automatically become the control node. The existing control node always remains as the control node unless it stops responding, at which point a new control node is elected.
5. In a "split-brain" scenario where multiple control nodes exist temporarily, the node with the highest priority retains the role, while the other nodes return to data node roles.

Note You can manually force a node to become the control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node.

High availability within the cluster

Clustering provides high availability by monitoring node and interface health and by replicating connection states between nodes.

Node health monitoring

Each node periodically sends a broadcast heartbeat packet over the cluster control link. If the control node does not receive any heartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining nodes.

Split brain scenario

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role.

Interface monitoring

Each node monitors the link status of all named hardware interfaces in use, and reports status changes to the control node.

All physical interfaces are monitored; only named interfaces can be monitored. You can optionally disable monitoring per interface.

A node is removed from the cluster if its monitored interfaces fail. The node is removed after 500 milliseconds.

Status after failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The Firewall Threat Defense automatically tries to rejoin the cluster, depending on the failure event.

Note When the Firewall Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoining the cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firewall Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firewall Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firewall Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from Cloud-Delivered Firewall Management Center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.

Data path connection state replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 55: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—

Traffic	State Support	Notes
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the cluster manages connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.
- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not

used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

Port Address Translation Connections

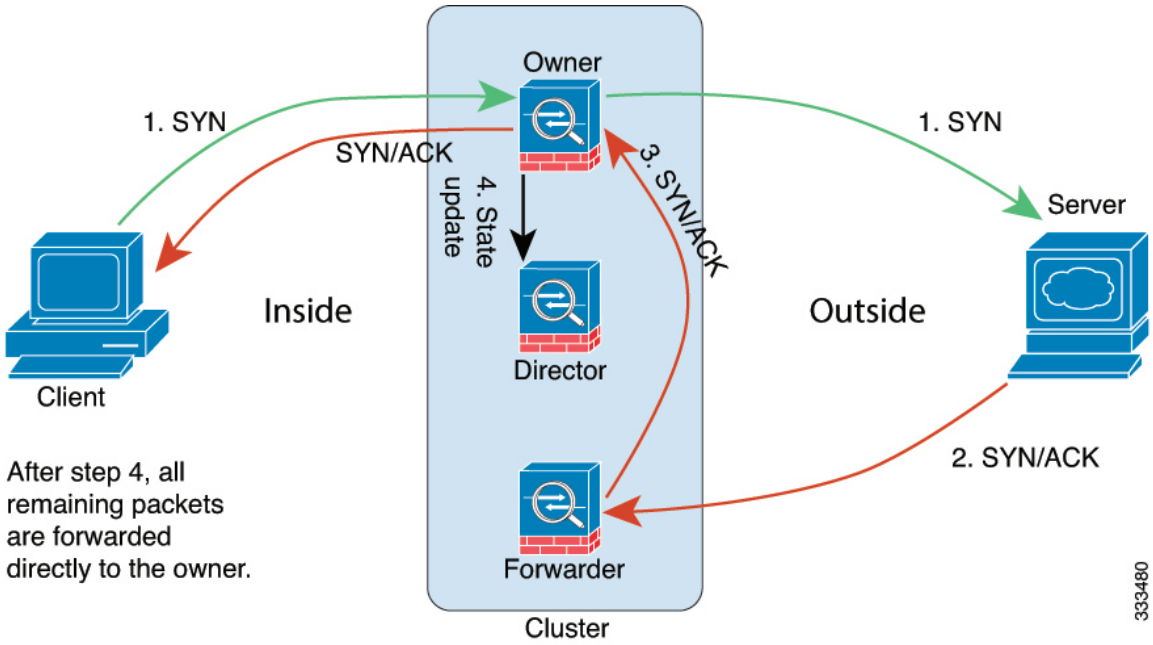
New connection ownership

When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. If a reverse flow arrives at a different node, it is redirected back to the original node.

Traffic redirection is not supported in this release. When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. All the subsequent packets for the same connection should arrive the same node. If any connection packets arrive at a different node, they will be dropped. If a reverse flow arrives at a different node, it will be dropped as well. For centralized features, if the connections do not arrive on the control node, they will be dropped.

Sample data flow for TCP

The following example shows the establishment of a new connection.

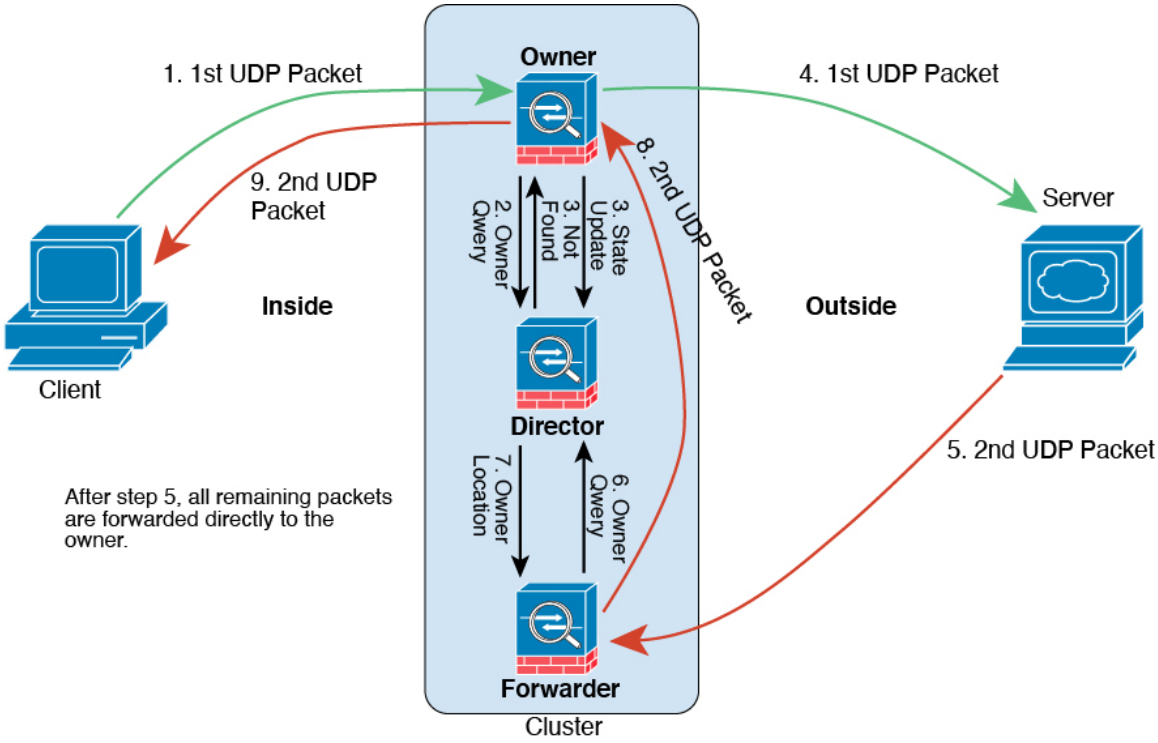


1. The SYN packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different Firewall Threat Defense (based on the load balancing method). This Firewall Threat Defense is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample data flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. *Figure 168: ICMP and UDP Data Flow*



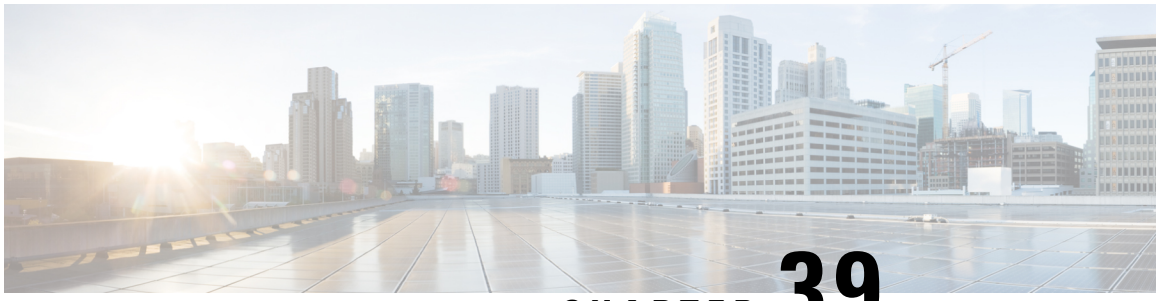
The first UDP packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for clustering: private cloud

This reference provides the feature history for clustering functionality in private cloud environments, including version information, platform compatibility, and detailed descriptions of enhancements and new capabilities.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
MTU ping test from data node on node join	20241030	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size twice the cluster control link MTU. Previously, only the control node sent a ping. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again. Added/modified commands: **show cluster history**. See: Clustering for the Secure Firewall 3100/4200, Clustering for Threat Defense Virtual in a Private Cloud, Clustering for Threat Defense Virtual in a Public Cloud, Clustering for the Firepower 4100/9300
Cluster health monitor settings	20221213	Any	You can now edit cluster health monitor settings. New/Modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard	20221213	Any	You can now view cluster health on the cluster health monitor dashboard. New/Modified screens: **System** > **Health** > **Monitor**
Clustering for the Firewall Threat Defense Virtual on VMware and KVM	Any	7.2.0	The Firewall Threat Defense Virtual supports Individual interface clustering for up to 4 nodes on VMware and KVM. New/Modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense Virtual on VMware and KVM



CHAPTER 39

Clustering: Public Cloud

Clustering lets you group multiple Firewall Threat Defense Virtuals together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. You can deploy Firewall Threat Defense Virtual clusters in a public cloud using the following public cloud platforms:

- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud Platform (GCP)

Currently, only routed firewall mode is supported.

Note Some features are not supported when using clustering. See Unsupported Features and Clustering.

- About Threat Defense Virtual clustering in the public cloud, on page 845
- Licenses for Threat Defense Virtual clustering, on page 848
- Requirements and prerequisites for Threat Defense Virtual clustering, on page 849
- Guidelines for Threat Defense Virtual clustering, on page 851
- Deploy the cluster in AWS, on page 852
- Deploy the cluster in Azure, on page 872
- Deploy the cluster in GCP, on page 891
- Add the cluster to the Management Center (manual deployment), on page 898
- Onboard a Secure Firewall Threat Defense Cluster, on page 904
- Edit cluster health monitor settings, on page 907
- Manage cluster nodes, on page 910
- Monitor cluster status, on page 913
- Upgrade the cluster, on page 919
- Reference for clustering, on page 919
- History for Threat Defense Virtual clustering in the public cloud, on page 932

About Threat Defense Virtual clustering in the public cloud

This section describes the clustering architecture and how it works.

How the cluster fits into your network

The cluster consists of multiple firewalls acting as a single device. To act as a cluster, the firewalls need the following infrastructure:

- Isolated network for intra-cluster communication, known as the cluster control link, using VXLAN interfaces. VXLANs, which act as Layer 2 virtual networks over Layer 3 physical networks, let the Firewall Threat Defense Virtual send broadcast or multicast messages over the cluster control link.
- Load Balancer(s)—For external load balancing, you have the following options depending on your public cloud:
  - AWS Gateway Load Balancer

    The AWS Gateway Load Balancer combines a transparent network gateway and a load balancer that distributes traffic and scales virtual appliances on demand. The Firewall Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint) using a Geneve interface single-arm proxy.
  - Azure Gateway Load Balancer

    In an Azure service chain, Firewall Threat Defense Virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The Firewall Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy.
  - Native GCP load balancers, internal and external
  - Equal-Cost Multi-Path Routing (ECMP) using inside and outside routers such as Cisco Cloud Services Router

    ECMP routing can forward packets over multiple “best paths” that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses or source and destination ports, or both, can be used to send a packet to one of the next hops. If you use static routes for ECMP routing, then the Firewall Threat Defense failure can cause problems; the route continues to be used, and traffic to the failed Firewall Threat Defense will be lost. If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each Firewall Threat Defense to participate in dynamic routing.



Note Layer 2 Spanned EtherChannels are not supported for load balancing.

Individual interfaces

You can configure cluster interfaces as individual interfaces.

Individual interfaces are normal routed interfaces, each with their own local IP address. The IP address for the interface will be configured automatically via DHCP. Static IP configuration is not supported.

Control and data node roles

When you add a cluster in the Cloud-Delivered Firewall Management Center , you choose one firewall to be the control node, and all additional firewalls to be data nodes. When you first create the cluster, the control node you specify will become the control node simply because it is the first node added to the cluster. Later, if multiple cluster nodes come online at the same time, the control node is determined by the priority setting; the priority is set between 1 and 100, where 1 is the highest priority.

When you create the cluster and deploy the cluster configuration, the Cloud-Delivered Firewall Management Center deploys a special bootstrap configuration to each node that includes fundamental settings such as the cluster control link configuration. The bootstrap configuration lets each node join the cluster. Most bootstrap settings are defined in the cluster wizard; however, you define the cluster control link interface hardware configuration (for example, creating an EtherChannel or setting the Ethernet speed) on the control node before you create the cluster. The cluster control link settings are copied to the bootstrap configuration for each node.

All nodes in the cluster share the same policy configuration. The node that you initially specify as the control node will overwrite the configuration on the data nodes when they join the cluster, so you only need to perform initial configuration on the control node before you form the cluster. After you create the cluster, all configuration changes are made at the cluster level and are shared by all nodes.

Some features do not scale in a cluster, and the control node handles all traffic for those features.

Cluster Control Link

Each node must dedicate one interface as a VXLAN (VTEP) interface for the cluster control link. For more information about VXLAN, see Configure VXLAN interfaces, on page 1071 .

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

VTEP Source Interface

The VTEP source interface is a regular Firewall Threat Defense Virtual interface with which you plan to associate the VNI interface. You can configure one VTEP source interface to act as the cluster control link. The source interface is reserved for cluster control link use only. Each VTEP source interface has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the cluster control link interfaces.

VNI Interface

A VNI interface is similar to a VLAN interface. It is a virtual interface that keeps network traffic separated on a given physical interface by using tagging. You can only configure one VNI interface. Each VNI interface has an IP address on the same subnet.

Peer VTEPs

Unlike regular VXLAN for data interfaces, which allows a single VTEP peer, Firewall Threat Defense Virtual clustering allows you to configure multiple peers.

Cluster control link traffic overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

Configuration replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Management network

You must manage each node using the management interface. Management from a data interface is not supported with clustering.

Licenses for Threat Defense Virtual clustering

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.



Note APEX, PLUS, and VPN-Only licenses are not supported for clustering. Before you create a cluster, remove unsupported license assignments from the devices that will become cluster nodes.

Performance tier license requirements

Each Firewall Threat Defense Virtual cluster node requires the same performance tier license. We recommend using the same number of CPUs and memory for all members, or else performance will be limited on all nodes to match the least capable member. The throughput level will be replicated from the control node to each data node so they match.

When you add the control node to the Cloud-Delivered Firewall Management Center, you can specify the feature licenses you want to use for the cluster. You can modify licenses for the cluster in the **Devices** > **Device Management**, **Cluster** > **License** area.

Note If you add the cluster before the Cloud-Delivered Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Cloud-Delivered Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and prerequisites for Threat Defense Virtual clustering

This reference provides the requirements and prerequisites for Threat Defense Virtual clustering, including model specifications, user roles, hardware and software requirements, and MTU configuration.

Model requirements

- FTDv5, FTDv10, FTDv20, FTDv30, FTDv50, FTDv100

Note FTDv5 and FTDv10 do not support Amazon Web Services (AWS) Gateway Load Balancer (GWLB) and Azure GWLB.

- Supported public cloud services:
  - Amazon Web Services (AWS)
  - Microsoft Azure
  - Google Cloud Platform (GCP)
- Maximum of 16 nodes

See also the general requirements for the Firewall Threat Defense Virtual in the Secure Firewall Threat Defense Virtual getting started guides.

User roles

- Admin
- Access Admin
- Network Admin

Hardware and software requirements

All units in a cluster:

- Must be in the same performance tier. We recommend using the same number of CPUs and memory for all nodes, or else performance will be limited on all nodes to match the least capable node.

- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported.
- The Cloud-Delivered Firewall Management Center access must be from the Management interface. Data interface management is not supported.
- Multiple Availability Zone deployment is supported from Firewall Threat Defense Virtual version 7.6.0 and later. The earlier versions supports only single Availability Zone deployment.
- Cluster requires a minimum of *four* interfaces: Management, Diagnostics, Data (Geneve), and CCL (Cluster Control Link).
- Cluster control link interfaces of all units must be in the same subnet.

MTU

Make sure the ports connected to the cluster control link have the correct (higher) MTU configured. If there is an MTU mismatch, the cluster formation will fail. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the initial ping fails, the node tries a ping using a smaller packet size (the MTU divided by 2, then by 4, then by 8) until a ping succeeds. A notification is generated so you can fix the MTU mismatch on connecting switches and try again.

The cluster control link MTU should be 154 bytes higher than the data interfaces. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead (100 bytes) plus VXLAN overhead (54 bytes).

For AWS with GWLB, the data interface uses Geneve encapsulation. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. You should set the source interface MTU to be the network MTU + 326 bytes. So for the standard 1500 MTU network path, the source interface MTU should be 1826, and the cluster control link MTU should be +154, 1980.

For Azure with GWLB, the data interface uses VXLAN encapsulation. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. You should set the cluster control link MTU to be the source interface MTU + 80 bytes.

The table shows the default values for the cluster control link MTU and the data interface MTU.

Note

- We do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.
- For Threat Defense Virtual clusters on Azure, modifying the cluster control link (CCL) MTU via FlexConfig in Management Center is not recommended, as it may cause cluster stability issues.

Table 56: Default MTU

Public Cloud	Cluster Control Link MTU	Data Interface MTU
AWS with GWLB	1980	1826
AWS	1654	1500
Azure with GWLB	1454	1374

Public Cloud	Cluster Control Link MTU	Data Interface MTU
Azure	1454	1300
GCP	1554	1400

Guidelines for Threat Defense Virtual clustering

High availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Multi-zone clustering

Multi-zone clustering is supported for up to three zones.

Additional guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Firewall Threat Defense or the switch, adding an additional switch to form a VSS or vPC or VNet) you should disable the health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.
- When adding a node to an existing cluster, or when reloading a node, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- Do not power off a node without first disabling clustering on the node.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new node. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.
- Dynamic scaling is supported from Secure Firewall version 7.3.
- Stateful Target Failover is not supported when you deploy the cluster on AWS.
- Perform a global deployment after the completion of each maintenance window.
- Ensure that you do not remove more than one device at a time from the auto scale group (AWS) / instance group (GCP) / scale set (Azure). We also recommend that you run the **cluster disable** command on the device before removing the device from the auto scale group (AWS) / instance group (GCP) / scale set (Azure).
- If you want to disable data nodes and the control node in a cluster, we recommend that you disable the data nodes before disabling the control node. If a control node is disabled while there are other data nodes

in the cluster, one of the data nodes has to be promoted to be the control node. Note that the role change could disturb the cluster.

- In the customized day 0 configuration scripts given in this guide, you can change the IP addresses as per your requirement, provide custom interface names, and change the sequence of the CCL-Link interface.
- If you experience CCL instability issues, such as intermittent ping failures, after deploying a Threat Defense Virtual cluster on a cloud platform, we recommend that you address the reasons that are causing CCL instability. Also, you can increase the hold time as a temporary workaround to mitigate CCL instability issues to a certain extent. For more information on how to change the hold time, see Edit Cluster Health Monitor Settings.
- When you are configuring your security firewall rule or security group for the Management Center virtual, you must include both Private and Public IP addresses of the Firewall Threat Defense Virtual in the Source IP address range. Also, ensure to specify the Private and Public IP addresses of the Firewall Management Center Virtual in the security firewall rule or security group of the Firewall Threat Defense Virtual. This is important to ensure proper registration of nodes during clustering deployment.

Defaults for clustering

- The cLACP system ID is auto-generated, and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

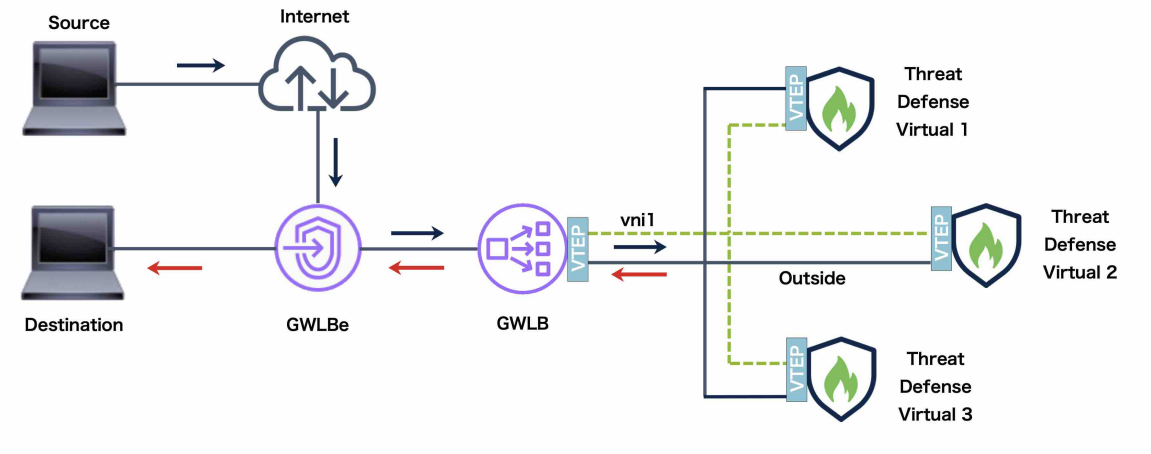
Deploy the cluster in AWS

You can deploy a cluster in AWS manually or by using CloudFormation templates to deploy a stack. You can use the cluster with AWS Gateway Load Balancer, or with a non-native load balancer such as the Cisco Cloud Services Router.

AWS gateway load balancer and Geneve single-arm proxy

This use case is the only currently supported use case for Geneve interfaces. The AWS Gateway Load Balancer combines a transparent network gateway and a load balancer that distributes traffic and scales virtual appliances on demand. The Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint). The following figure shows traffic forwarded to the Gateway Load Balancer from the Gateway Load Balancer endpoint. The Gateway Load Balancer balances traffic among multiple Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer (U-turn traffic). The Gateway Load Balancer then sends the traffic back to the Gateway Load Balancer endpoint and to the destination.

Figure 169: Geneve Single-Arm Proxy



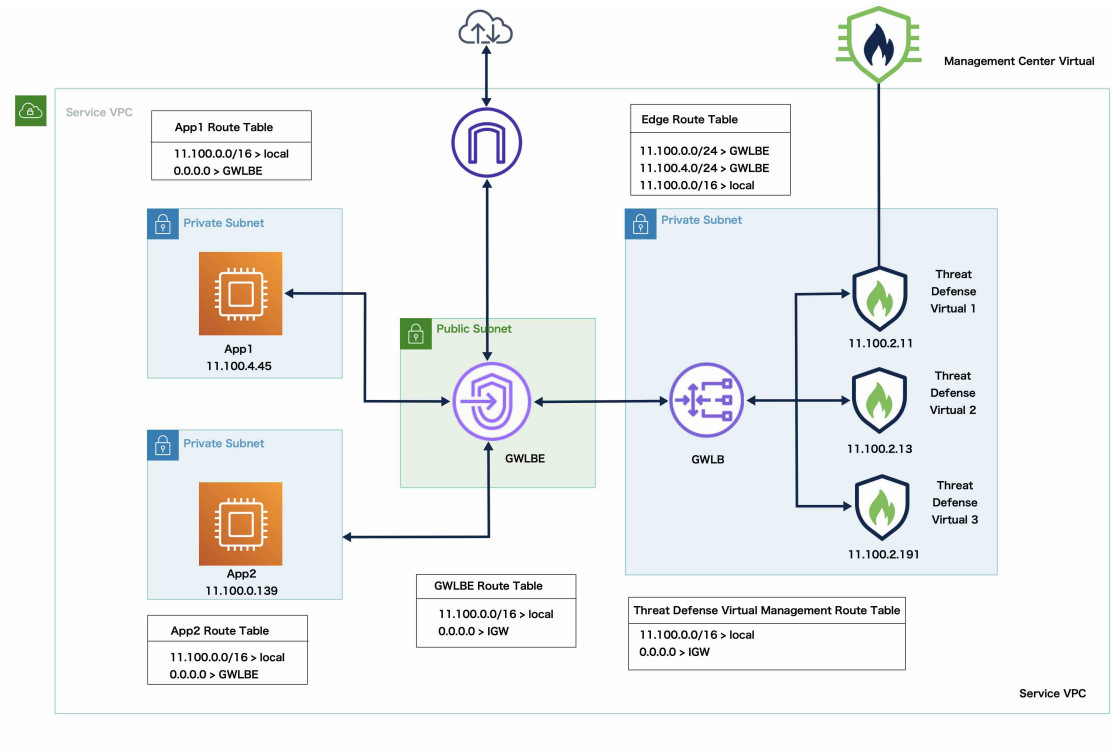
Note Transport Layer Security (TLS) Server Identity Discovery is not supported with Geneve single-arm setup on AWS.

Sample topology

The sample topology provided below

- depicts both inbound and outbound traffic flow through security appliances,
- includes three Threat Defense Virtual instances in a cluster connected to a GWLB, and
- uses a Management Center Virtual instance to manage the cluster.

Traffic flow



Inbound traffic from the internet goes to the GWLB endpoint which then transmits the traffic to the GWLB. Traffic is then forwarded to the Threat Defense Virtual cluster. After the traffic has been inspected by a Threat Defense Virtual instance in the cluster, it is forwarded to the application VM, App1 /App2.

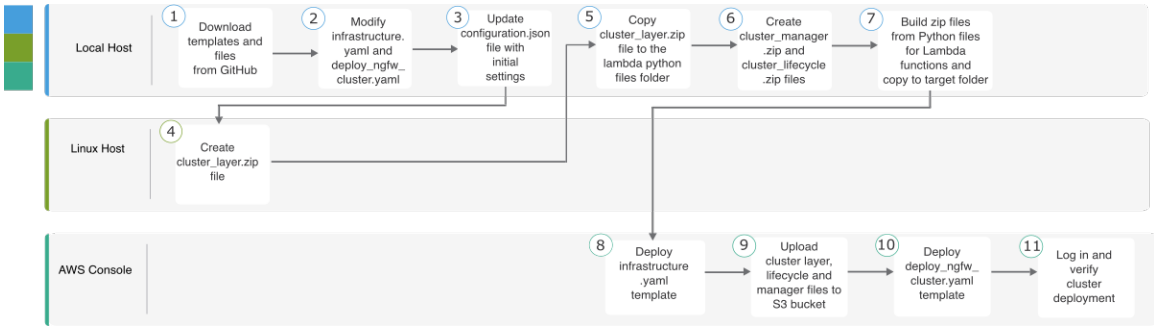
Outbound traffic from App1/App2 is transmitted to the GWLB endpoint which then sends it out to the internet.

End-to-end process for deploying Threat Defense Virtual cluster on AWS

Provides a workflow to help users successfully deploy and verify a Threat Defense Virtual cluster on AWS using either template-based or manual deployment processes.

Template-based Deployment

The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster on AWS.



	Workspace	Steps
1	Local Host	Clone the repository from GitHub
2	Local Host	Modify *infrastructure.yaml* and *deploy_ngfw_cluster.yaml* templates.
3	Local Host	Update the *Configuration.json* file with FMC object names.
4	Linux Host	Create *cluster_layer.zip* file.
5	Local Host	Copy *cluster_layer.zip* file to the Lambda python files folder.
6	Local Host	Create *cluster_manager.zip*, *custom_metrics_publisher.zip*, and *cluster_lifecycle.zip* files.
7	Local Host	Build zip files from Python files for Lambda functions and copy to target folder.
8	AWS Console	Deploy *infrastructure.yaml* template.
9	AWS Console	Upload *cluster_layer.zip*, *cluster_lifecycle.zip*, *custom_metrics_publisher.zip*, and *cluster_manager.zip* to the S3 bucket.
10	AWS Console	Deploy *deploy_ngfw_cluster.yaml* template.
11	AWS Console	Log in and verify cluster deployment.

Manual Deployment

The following flowchart illustrates the workflow for manual deployment of the Threat Defense Virtual cluster on AWS.



	Workspace	Steps
1	Local Host	Create day 0 configuration script.
2	AWS Console	Deploy Threat Defense Virtual instance.
3	AWS Console	Attach interfaces to instance.
4	AWS Console	Verify if nodes have joined cluster.
5	AWS Console	Create target group and GWLB; attach target group to the GWLB.
6	AWS Console	Register instances with the target group using data interface IP.
7	Management Center	Register control node.

Templates

The templates are available in GitHub. The parameter values are self-explanatory with the parameter names, default values, allowed values, and description, given in the template.

- infrastructure.yaml – Template for infrastructure deployment.
- deploy_ngfw_cluster.yaml – Template for cluster deployment.

Note Ensure that you check the list of supported AWS instance types before deploying cluster nodes. This list is found in the *deploy_ngfw_cluster.yaml* template, under allowed values for the parameter InstanceType.

Deploy the Stack in AWS Using a CloudFormation template

Deploy the stack in AWS using the customized CloudFormation template.

Before you begin

- You need an Amazon Linux virtual machine with Python 3.
- To allow the cluster to auto-register with the Cloud-Delivered Firewall Management Center , you need to create two users with administrative privileges on the Cloud-Delivered Firewall Management Center that can use the REST API. See the *Cisco Secure Firewall Management Center Administration Guide* .
- Add an access policy in the Cloud-Delivered Firewall Management Center that matches the name of the policy that you specified in `Configuration.json` .

Procedure

Step 1 Prepare the template.

a) Clone the GitHub repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/aws .

b) Modify **infrastructure.yaml** and **deploy_ngfw_cluster.yaml** with the required parameters.

c) Modify **cluster/aws/lambda-python-files/Configuration.json** with initial settings.

For example:

```
{
  "licenseCaps": ["BASE", "MALWARE", "THREAT"],
  "performanceTier": "FTDv50",
  "fmcIpforDeviceReg": "DONTRESOLVE",
  "RegistrationId": "cisco",
  "NatId": "cisco",
  "fmcAccessPolicyName": "AWS-ACL"
}
```

- Keep the fmcIpforDeviceReg setting as DONTRESOLVE.
- The fmcAccessPolicyName needs to match an access policy on the Cloud-Delivered Firewall Management Center .

Note
FTDv5 and FTDv10 tiers are not supported.

d) Create a file named **cluster_layer.zip** to provide essential Python libraries to Lambda functions.

We recommend using Amazon Linux with Python 3.9 installed to create the cluster_layer.zip file.

Note
If you need an Amazon Linux environment, you can create an EC2 instance using Amazon Linux 2023 AMI or use AWS Cloudshell, which runs the latest version of Amazon Linux.

For creating the cluster-layer.zip file, you need to first create requirements.txt file that consists of the python library package details and then run the shell script.

1. Create the requirements.txt file by specifying the python package details.

   The following is the sample package details that you provide in the requirements.txt file:

   ```
   $ cat requirements.txt
   pycryptodome
   paramiko
   requests
   scp
   jsonschema
   cffi
   zipp
   importlib-metadata
   ```

2. Run the following shell script to create cluster_layer.zip file.

   ```
   $ pip3 install --platform manylinux2014_x86_64
   --target=./python/lib/python3.9/site-packages
   --implementation cp --python-version 3.9 --only-binary=:all:
   ```

```
--upgrade -r requirements.txt
$ zip -r cluster_layer.zip ./python
```

Note
If you encounter a dependency conflict error during installation, such as urllib3 or cryptography, it is recommended that you include the conflicting packages along with their recommended versions in the requirements.txt file. After that, you can run the installation again to resolve the conflict.

e) Copy the resulting **cluster_layer.zip** file to the lambda python files folder - `cluster/aws/lambda-python-files` .

f) Create the **cluster_layer.zip** , **custom_metrics_publisher.zip** , **cluster_manger.zip** and **lifecycle_ftdv.zip** files.

A make.py file can be found in the cloned repository (`cluster/aws/make.py`). This will zip the python files into a Zip file and copy to a target folder.

python3 make.py build

Note
If you are using a private IP address for the Management Center Virtual registration, then make sure that you set USE_PUBLIC_IP_FOR_FMC_CONN to False in the `cisco-ftdv/cluster/aws/lambda-python-files/constant.py` file.

Step 2 Deploy **infrastructure.yaml** and note the output values for cluster deployment. Before deploying the infrastructure stack, it is important to identify the AWS region and the availability zones that will be used. Each AWS region has a different set of availability zones and VPC infrastructure. Therefore, it is essential to select the correct region and availability zones for your deployment.

a) On the AWS Console, go to **CloudFormation** and click **Create stack** ; select **With new resources(standard)** .

b) Select **Upload a template file** , click **Choose file** , and select **infrastructure.yaml** from the target folder.

c) Click **Next** and provide the required information.

e

Parameter	Allowed Values/Type	Description
ClusterName	String	Enter unique Cluster name.
ClusterNumber	Number	Enter unique Cluster number.
VpcCidr	String	Enter the CIDR block for a new VPC
NoOfAZs	Number	Select 2 or 3 Availability Zones (AZs) for releases 7.6.0 and above; for lower releases, Select 1AZ. Management, Inside, Outside and CCL subnets will be distributed across the chosen AZs accordingly.
ListOfAZs	List	Select Availability Zones (Count should match with Number of Availability Zones)
MgmtSubnetNames	CommaDelimitedList	Management subnets name (With Internet GW as Route)

Parameter	Allowed Values/Type	Description
MgmtSubnetCidrs	CommaDelimitedList	Management subnets Cidr list
InsideSubnetNames	CommaDelimitedList	Inside subnets name (With Private Route)
InsideSubnetCidrs	CommaDelimitedList	Inside subnets Cidr list
CCLSubnetNames	CommaDelimitedList	Enter CCL subnet name
CCLSubnetCidrs	CommaDelimitedList	Enter CCL subnet CIDR
LambdaAZs	List	Select 2 Availability Zones for Lambda
LambdaSubnetNames	CommaDelimitedList	Enter Lambda Subnets name (With NAT GW as Route), for Lambda Functions
LambdaSubnetCidrs	CommaDelimitedList	Enter Lambda Subnet CIDRs

d) Enter a unique **Cluster Name** and **Cluster Number** for the cluster.
e) Select the availability zone from the **Availability Zone** list. This field lists only availability zones based on the AWS region that you select for deploying the infrastructure stack using the ClusterFormation template.
f) Click **Next** , then **Create stack** .
g) After the deployment is complete, go to **Outputs** and note the S3 **BucketName** .

Figure 170: Output of infrastructure.yaml

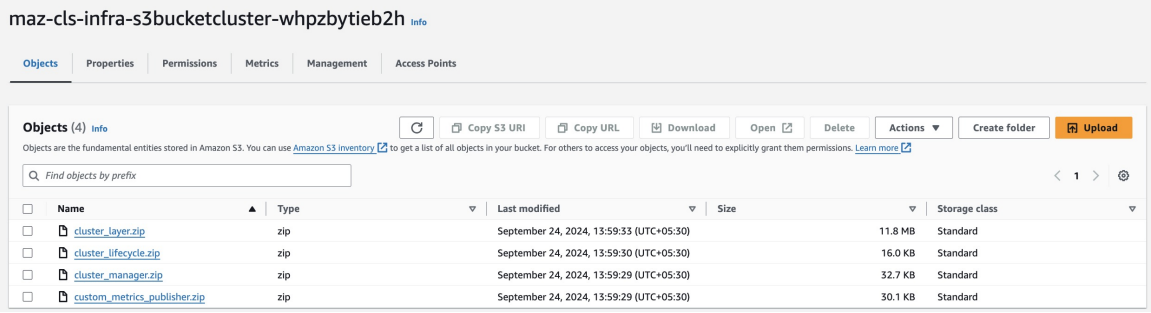
Outputs (13)

Search outputs

Key	Value	Description	Export name
BucketName	maz-cls-infra-s3bucketcluster-whpzbytieb2h	Name of the Amazon S3 bucket	-
BucketUrl	http://maz-cls-infra-s3bucketcluster-whpzbytieb2h.s3-website-us-east-1.amazonaws.com	URL of S3 Bucket Static Website	-
CCLSubnetIds	subnet-0bc04e2cc9e53e5c0,subnet-0d7d046a0fca25615,subnet-03ef42bf527551569	List of CCL subnet IDs (comma seperated)	-
EIPforNATgw	3.218.44.132	EIP reserved for NAT GW	-
FmcInstanceSGID	sg-076880aa64df2db5c	Security Group ID for FMC if user would like to launch in this VPC itself	-
InInterfaceSGId	sg-06ed933d6624fe51b	Security Group ID for Inside Interfaces	-
InsideSubnetIds	subnet-03d12cab8ee0eafff,subnet-0be9158b0970aebab,subnet-0b53c96fceb7c1f4d	List of Inside subnet IDs (comma seperated)	-
InstanceSGId	sg-0680b74be473186aa	Security Group ID for Instances Management Interface	-
LambdaSecurityGroupId	sg-057da2a9954e0d204	Security Group ID for Lambda Functions	-
LambdaSubnetIds	subnet-03439803d989e6bdf,subnet-087488a9d6ffc95cd	List of lambda subnet IDs (comma seperated)	-
ListOfAZs	us-east-1a,us-east-1b,us-east-1c	Availability zones for NGFWv instances	-
MgmtSubnetIds	subnet-06f0bbbd3f207a504,subnet-0c339dc43688cddc9,subnet-0a67629632a655de7	List of Mangement subnet IDs (comma seperated)	-
VpcName	vpc-09c2b0ad995e2fb24	Name of the VPC created	-

Step 3 Upload **cluster_layer.zip** , **cluster_manager.zip** , **custom_metrics_publisher.zip** , and **cluster_lifecycle.zip** to the S3 bucket created by **infrastructure.yaml** .

Figure 171: S3 Bucket



Note
Make sure that the Elastic IP address of the Lambda NAT Gateway is added to the security group associated with the Management Center Virtual.

Step 4 Deploy **deploy_ngfw_cluster.yaml** .

a) Go to **CloudFormation** and click on **Create stack** ; select **With new resources(standard)** .
b) Select **Upload a template file** , click **Choose file** , and select **deploy_ngfw_cluster.yaml** from the target folder.
c) Click **Next** and provide the required information.
d) Provide the following cluster and infrastructure configuration information.

Parameter	Allowed Values/Type	Description
Cluster Configuration		
ClusterGrpNamePrefix	String	This is the cluster name Prefix. The cluster number will be added as a suffix.
ClusterNumber	String	This is the cluster number. This will be suffixed to the cluster name (msa-ftdv-infra). For example, if this value is 1 , the group name will be msa-ftdv-infra-1. It should be at least 1 digit, but not more than 3 digits. Default: 1.
ClusterSize	Numbers	This is the total number of Firewall Threat Defense Virtual nodes in a cluster. Minimum: 1 Maximum:16
Infrastructure Details		
NoOfAZs	String	This is the total number of availability zones into which Firewall Threat Defense Virtual is deployed. (The number of availability zones varies from a Minimum 1 to Maximum 3 depending on a region). The subnet will be created in these availability zones.

Parameter	Allowed Values/Type	Description
		The availability zones available in this list is based on the region selected for deploying the cluster. **Note** Management, Inside, and Cluster Control Link (CCL) subnets are created across three availability zones based on this parameters.
AZ	String	The availability zone list is based on the region you plan to deploy. In Availability Zone list, select the valid availability zone (1 availability zone or 2 availability zones or 3 availability zones). Count should match with the value of Number of Availability Zones parameter.
NotifyEmailID	String	Email address to which cluster events email will be sent. You must accept a subscription email request to receive this email notification. For example, admin@company.com.
VpcId	String	The VPC ID for the cluster group. Type: AWS::EC2::VPC::Id
S3BktName	String	The S3 Bucket that contains the uploaded Lambda zip files. You must specify correct bucket name.
MgmtSubnetIds	List	Enter only one subnet per availability zone. If you select multiple subnets from a same availability zone, then selecting an incorrect subnet may cause issues while deploying the Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::Subnet::Id>
InsideSubnetIds	List	Enter at least one subnet per availability zone. If multiple subnets from the same Availability Zone are selected, then selecting an incorrect subnet may cause issues while deploying the Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::Subnet::Id>
LambdaSubnets	List	Enter at least two subnets for the Lambda functions. The two subnets you enter must have a NAT gateway to enable the Lambda functions to communicate with AWS services, which are public DNS. Type: List<AWS::EC2::Subnet::Id>

Parameter	Allowed Values/Type	Description
CCLSubnetIds	String	Enter at least one subnet per availability zone. If multiple subnets from the same Availability Zone are selected, then selecting an incorrect subnet may cause issues while deploying the Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::Subnet::Id>
CCLSubnetRanges	String	Enter IP addresses range of CCL subnets for different availability zones. Exclude first 4 reserved IP addresses. IP address pool for Cluster Control Link (CCL). IP address is allocated to the CCL interfaces of the Firewall Threat Defense Virtual instance from CCL IP address pool.
MgmtInterfaceSG	List	Select security group ID for the Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::SecurityGroup::Id>
InsideInterfaceSG	List	Select security group ID for the inside interface of Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::SecurityGroup::Id>
LambdaSG	List	Select a security group for the Lambda functions. Ensure outbound connections are set to ANYWHERE . Type: List<AWS::EC2::SecurityGroup::Id>
CCLInterfaceSG	List	Select a security group ID for CCL interface of the Firewall Threat Defense Virtual instances.
GWLB Configuration		
DeployGWLBE	String	Click Yes to deploy the GWLB endpoint. By default, the value is set to No .
VpcIdLBE	String	Enter VPC to deploy Gateway Load Balancer Endpoint. **Note** Do not enter any value in this field if you are not deploying the GWLB endpoint.
GWLBESubnetId	String	Enter only one subnet ID. **Note** Do not enter any value in this field if you are not deploying the GWLB endpoint.

Parameter	Allowed Values/Type	Description
		Ensure that the subnet belongs to the correct VPC, and the availability zones that you have specified.
TargetFailover	String	Enable Target Failover support when a target fails or deregisters. (By default, the value of this parameter is set to rebalance). • no_rebalance : Directs existing flows to failed targets and new flows to healthy targets, ensuring backward compatibility. • rebalance : Redistributes existing flows while ensuring that new flows go to healthy targets. rebalance is supported from Firewall Threat Defense Virtual Version 7.4.1 and later.
TgHealthPort	String	Enter Health Check Port for GWLB. **Note** By default, this port must not be used for traffic. Ensure the value you provide is a valid TCP port. Default: 8080
Cisco NGFWv Instance Configuration		
InstanceType	String	Cisco Firewall Threat Defense Virtual EC2 instance type. Ensure that the AWS Region supports Instance Type you select. By default, **c5.xlarge** is selected.
LicenseType	String	Choose Cisco Firewall Threat Defense Virtual EC2 instance license type. Ensure that the AMI ID that you enter in **AMI-ID** parameter is of the same licensing type. By default, **BYOL** is selected.
AssignPublicIP	String	Set the value as **true** to assign a public IP address for Firewall Threat Defense Virtual from the AWS IP address pool.
AmiID	String	Choose the correct AMI ID as per the region, version, and license type (BYOL or PAYG). Firewall Threat Defense Virtual 7.2 and later support clustering, and Firewall Threat Defense Virtual Version 7.6 and later support the autoscaling and multiple availability zone enhancements. Type: AWS::EC2::Image::Id

Parameter	Allowed Values/Type	Description
ngfwPassword	String	Firewall Threat Defense Virtual instance password. All Firewall Threat Defense Virtual instances come up with a default password, which is in the Userdata field of the Launch Template (Cluster Group). The password is activated after Firewall Threat Defense Virtual is accessible. Minimum length must be 8 characters. The password can either be a plain text password or a KMS encrypted password.
KmsArn	String	Enter ARN of an existing KMS (AWS KMS key to encrypt at rest). If you specify a value in this field, then the Firewall Threat Defense Virtual instance's admin password must be an encrypted password. Example of generating an encrypted password: " aws kms encrypt --key-id <KMS ARN> --plaintext <password> " The password encryption must be done using only the specified ARN.
FMC Automation Configuration		
fmcDeviceGrpName	String	Enter a unique name for the cluster group in management center.
fmcPublishMetrics	String	Select **`true`** to create a Lambda Function to poll management center and publish specific device group metrics to AWS CloudWatch. Allowed values: • true • false By default, the value is set to true .
fmcMetricsUsername	String	Enter a unique internal user name for polling memory metrics from management center. The user must have privileges of Network Admin and Maintenance User or more .
fmcMetricsPassword	String	Enter the password. If you have mentioned KMS Master Key ARN parameter, ensure to provide an encrypted password.

Parameter	Allowed Values/Type	Description
		Ensure to enter the correct password because entering incorrect password may result in failure of metrics collection.
fmcServer	String	The IP address can be an external IP address or the IP address reachable in Firewall Threat Defense Virtual management subnet in the VPC. Minimum length: 7 Maximum length:15
fmcOperationsUsername	String	Provide a unique internal user name for Firewall Management Center Virtual for CloudWatch. The user must have Administrator privileges.
fmcOperationsPassword	String	Enter the password. If you have mentioned KMS Master Key ARN parameter, ensure to provide an encrypted password.
Scaling Configuration		
CpuThresholds	CommaDelimitedList	(Optional) Specifying non-zero lower and upper thresholds will create scale policies. If (0,0) is selected, no CPU scaling alarm or policies will be created. Evaluation points and data points are at default or recommended values. By default, autoscaling is enabled in this template. Autoscale can be disabled after deployment.
MemoryThresholds	CommaDelimitedList	Specifying non-zero lower and upper threshold will create scale policies. If (0,0) is selected, no memory scaling alarm or policies will be created. Evaluation points and data points are at default or recommended values.

e) Click **Next** , then **Create stack** .

The Lambda functions manage the rest of the process, and the Firewall Threat Defense Virtual s will automatically register with the Cloud-Delivered Firewall Management Center .

Figure 172: Deployed Resources

The status changes from CREATE_IN_PROGRESS to CREATE COMPLETE, indicating successful deployment.

Step 5 Verify the cluster deployment by logging into any one of the nodes and using the **show cluster info** command.

Figure 173: Cluster Nodes

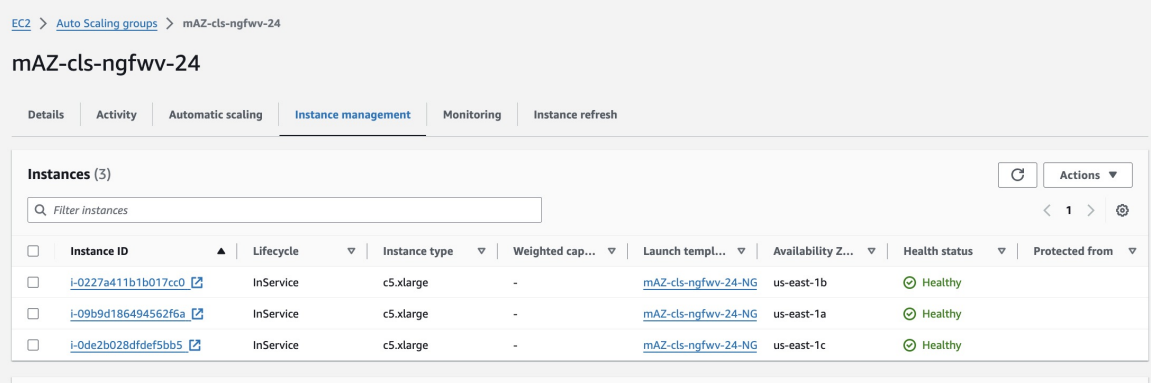
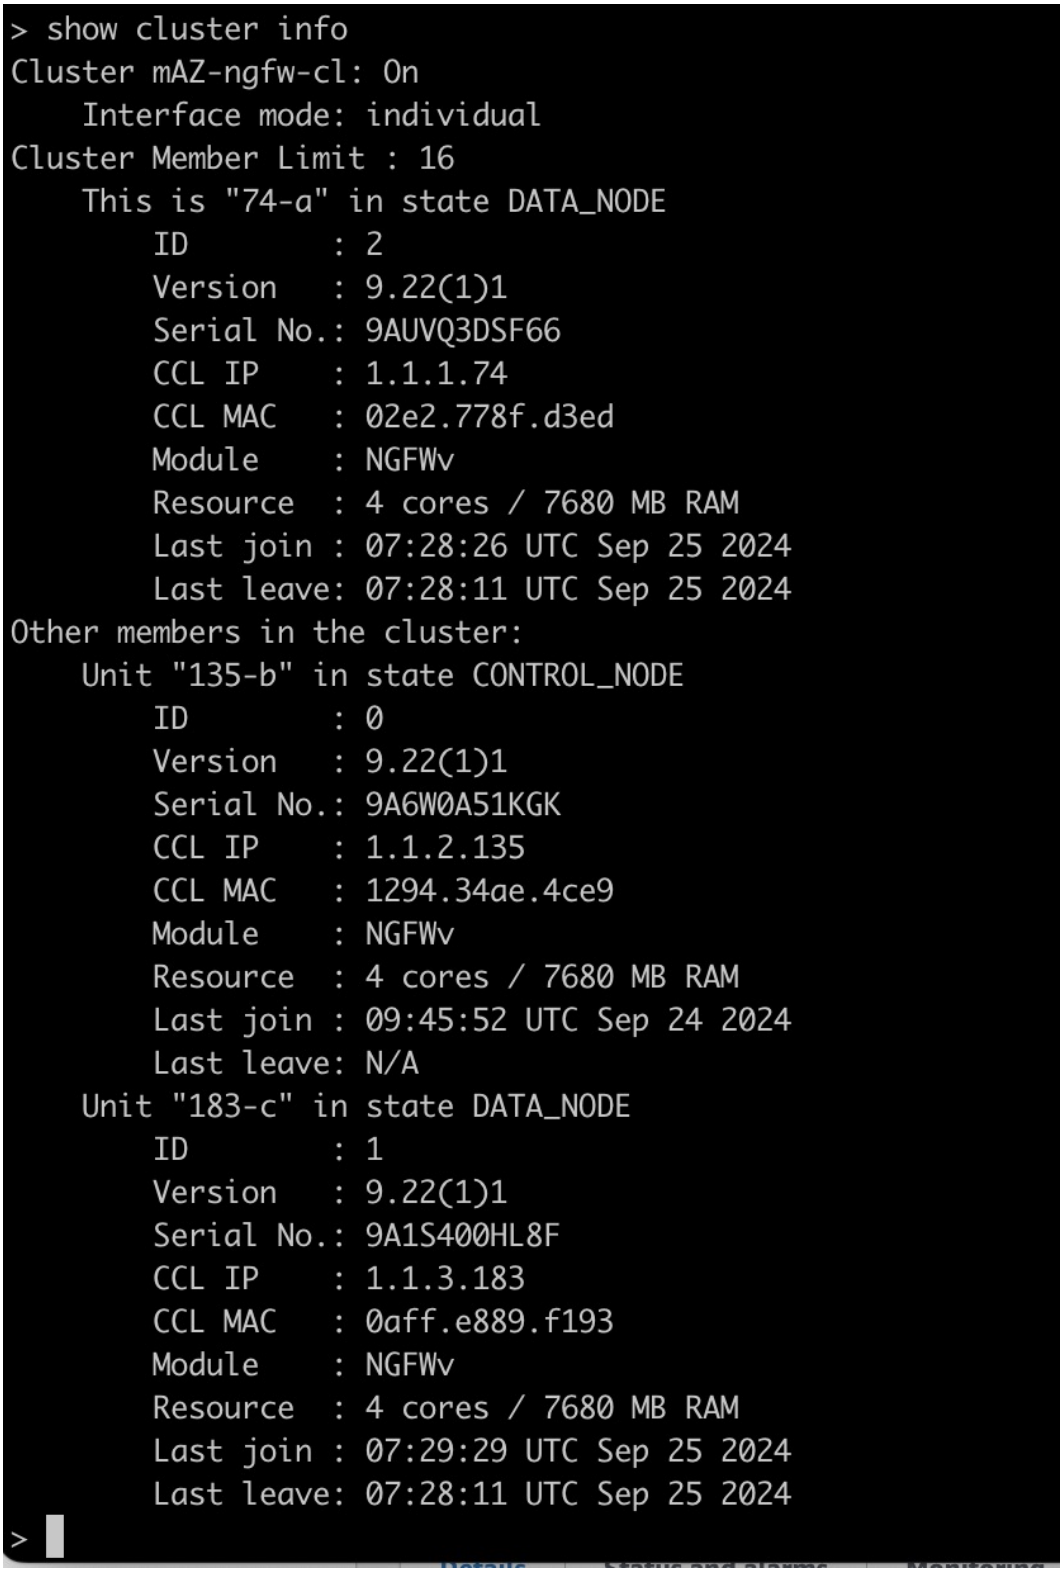


Figure 174: show cluster info



Manual cluster deployment in AWS

To deploy the cluster manually, prepare the day 0 configuration, deploy each node, and then add the control node to the Cloud-Delivered Firewall Management Center.

On version 7.6.4-69 and from 10.0.0, each cluster node requires a unique AWS instance tag "cluster-node-id" with a value ranging from 1 to 16 for internal cluster configuration. Please ensure that the tag is added before the device boots up.

For example: key "cluster-node-id" -> value "1"

Additionally, make sure that "Allow tags in instance metadata" is set to "Enabled."

Create the Day0 configuration for AWS

You can use either a fixed configuration or a customized configuration. We recommend using the fixed configuration.

Create the Day0 configuration with a fixed configuration for AWS

The fixed configuration will auto-generate the cluster bootstrap configuration.

Single availability zone - Day0 configuration with a fixed configuration for AWS

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "Cluster": {
        "CclSubnetRange": "ip_address_start ip_address_end",
        "ClusterGroupName": "cluster_name",
        [For Gateway Load Balancer] "Geneve": "{Yes | No}",
        [For Gateway Load Balancer] "HealthProbePort": "port"
    }
}
```

For example:

```
{
 "AdminPassword": "Sup3rnatural",
 "Hostname": "ciscoftdv",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "Cluster": {
              "CclSubnetRange": "10.5.90.4 10.5.90.30",
  "ClusterGroupName": "ftdv-cluster",
  "Geneve": "Yes",
  "HealthProbePort": "7777"
 }
}
```

Configure the cluster infrastructure subnet range

Use a cluster infrastructure subnet range to allocate IP addresses within a /27 CIDR block. This range provides the usable IP addresses for your cluster infrastructure.

Before you begin

- Review the table to determine valid start and end IP addresses for the /27 CIDR blocks you plan to use.
- The first four IP addresses in each subnet (including the network address) are reserved by AWS. Configure the **CclSubnetRange** to begin with the fifth IP address in the subnet.

The table shows the valid start and end IP addresses for each /27 CIDR allocation.

Table 57: CclSubnetRange mapping

CIDR	Start IP Address	End IP Address
x.y.z.0/27	x.y.z.4	x.y.z.30
x.y.z.32/27	x.y.z.36	x.y.z.62
x.y.z.64/27	x.y.z.68	x.y.z.94
x.y.z.96/27	x.y.z.100	x.y.z.126
x.y.z.128/27	x.y.z.132	x.y.z.158
x.y.z.160/27	x.y.z.164	x.y.z.190
x.y.z.192/27	x.y.z.196	x.y.z.222
x.y.z.224/27	x.y.z.228	x.y.z.254

Complete these steps to configure the cluster infrastructure subnet range.

Procedure

Step 1 Identify the /27 CIDR block assigned to your cluster infrastructure subnet.

Step 2 Find the matching CIDR in the mapping table.

Step 3 Replace "x.y.z" with the first three octets of your subnet. For example, 172.20.10 or 10.100.5.

Step 4 Configure the range by entering the start and end IP addresses from the mapping table for your subnet.

- Example 1:
  - CIDR: 172.20.10.64/27
  - Subnet range: 172.20.10.68 to 172.20.10.94
- Example 2:
  - CIDR: 10.100.5.128/27

- Subnet range: 10.100.5.132 to 10.100.5.158

Deploy cluster nodes

Deploy cluster nodes to form a functional cluster that can handle traffic distribution and provide high availability for network security services.

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Deploy the Threat Defense Virtual instance by using the cluster day 0 configuration with the required number of interfaces - four interfaces if you are using Gateway Load Balancer (GWLB), or five interfaces if you are using non-native load balancer. To do this, in the **Configure Instance Details** > **Advanced Details** section, paste the cluster day 0 configuration.

Note
Ensure that you attach interfaces to the instances in the order given below.

- AWS Gateway Load Balancer - four interfaces - management, diagnostic, inside, and cluster control link.
- Non-native load balancers - five interfaces - management, diagnostic, inside, outside, and cluster control link.

For more information on deploying Threat Defense Virtual on AWS, refer to Deploy the Threat Defense Virtual on AWS.

Step 2 Repeat Step 1 to deploy the required number of additional nodes.

Step 3 Use the **show cluster info** command on the Threat Defense Virtual console to verify if all nodes have successfully joined the cluster.

Step 4 Configure the AWS Gateway Load Balancer.

a) Create a target group and GWLB.

b) Attach the target group to the GWLB.

Note
Ensure that you configure the GWLB to use the correct security group, listener configuration, and health check settings.

c) Register the data interface (inside interface) with the Target Group using IP addresses.

For more information, refer to Create a Gateway Load Balancer.

Step 5 Add the control node to the Management Center. Refer to Onboard a Secure Firewall Threat Defense Cluster.

Deploy the cluster in Azure

You can use the cluster with the Azure Gateway Load Balancer (GWLB), or with a non-native load balancer. To deploy a cluster in Azure, use Azure Resource Manager (ARM) templates to deploy a Virtual Machine Scale Set.

Sample topology for GWLB-based cluster deployment

Sample topologies for deploying Gateway Load Balancer (GWLB) clusters are given below.

Figure 175: Inbound traffic use case and topology with GWLB

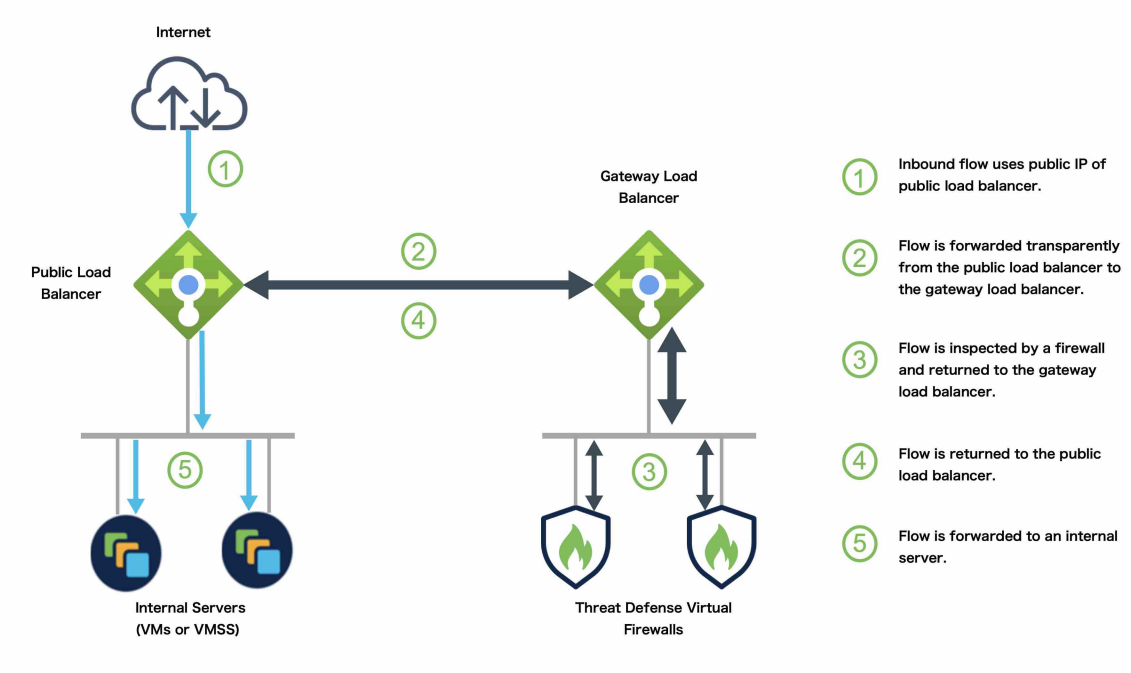
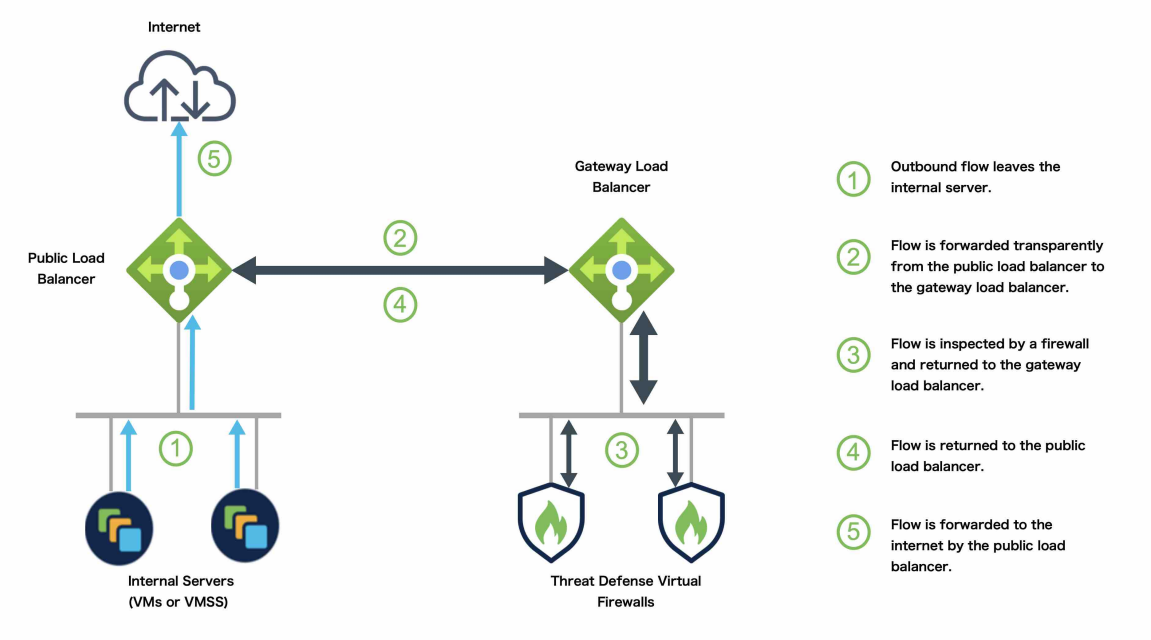


Figure 176: Outbound traffic use case and topology with GWLB

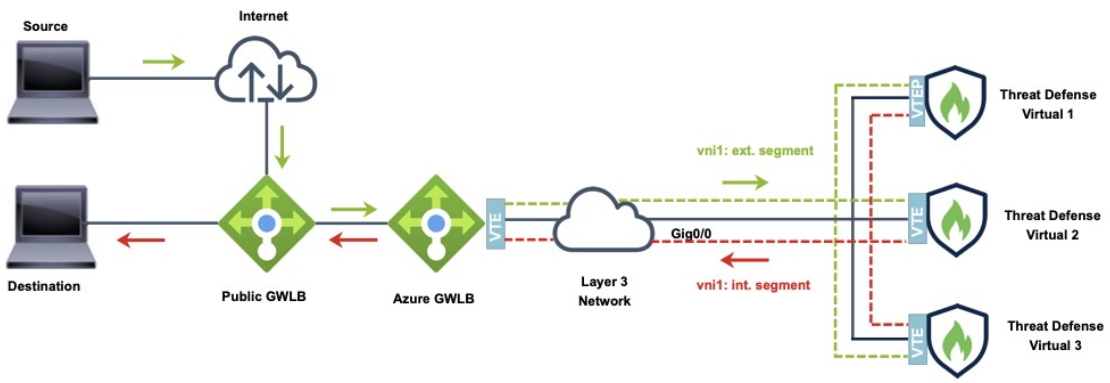


Azure gateway load balancer and paired proxy

In an Azure service chain, Threat Defense Virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy.

The figure shows traffic forwarded to the Azure Gateway Load Balancer from the Public Gateway Load Balancer on the external VXLAN segment. The Gateway Load Balancer balances traffic among multiple Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer on the internal VXLAN segment. The Azure Gateway Load Balancer then sends the traffic back to the Public Gateway Load Balancer and to the destination.

Figure 177: Azure gateway load balancer with paired proxy



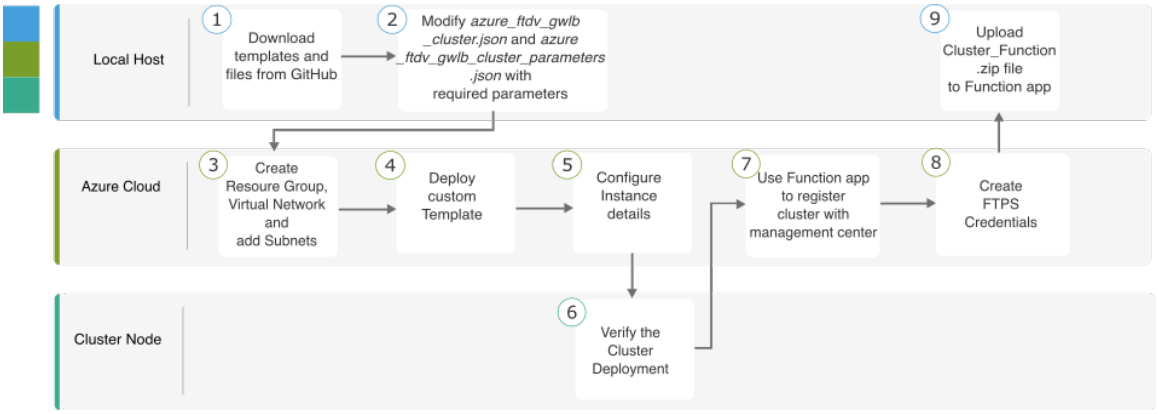
Traffic flow between GWLBe to GWLB (Geneve Single-Arm Proxy) in Azure

End-to-end process for deploying Threat Defense Virtual cluster in Azure with GWLB

Provides a workflow to help users successfully deploy and verify a Threat Defense Virtual cluster on Azure with GWLB using either template-based or manual deployment processes.

Template-based Deployment

The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster in Azure with GWLB.

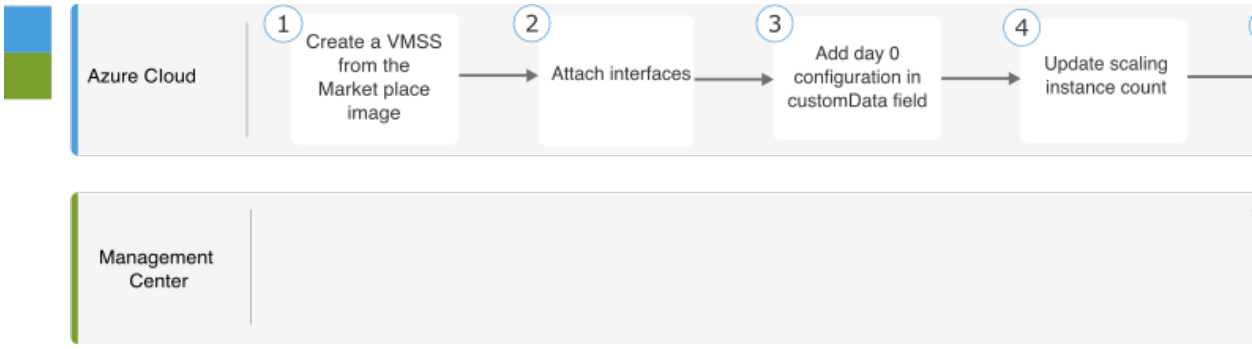


	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Modify *azure_ftdv_gwlb_cluster.json* and *azure_ftdv_gwlb_cluster_parameters.json* with the required parameters.

	Workspace	Steps
3	Azure Cloud	Create the resource group, virtual network, and subnets.
4	Azure Cloud	Deploy custom template.
5	Azure Cloud	Configure instance details.
6	Cluster Node	Verify cluster deployment.
7	Azure Cloud	Use the Function app to register the cluster with the Management Center.
8	Azure Cloud	Create FTPS credentials.
9	Local Host	Upload *Cluster_Function.zip* file to the Function app.

Manual Deployment

The following flowchart illustrates the workflow of manual deployment of Threat Defense Virtual cluster in Azure with GWLB.



	Workspace	Steps
1	Local Host	Create a VMSS from the Marketplace image.
2	Local Host	Attach interfaces.
3	Local Host	Add day 0 configuration in the customData field.
4	Local Host	Update scaling instance count.
5	Local Host	Configure GWLB.

	Workspace	Steps
6	Management Center	Add control node.

Templates

The templates are available in GitHub. The parameter values are self-explanatory with the parameter names and values given in the template.

Prerequisites for Azure GWLB cluster deployment

Complete these prerequisite tasks before deploying the Azure GWLB cluster.

- To allow the cluster to auto-register to the management center, create a user with Network Admin & Maintenance User privileges on the management center. Users with these privileges can use REST API. Refer to the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the management center that matches the name of the policy that you will specify during template deployment.
- Ensure that the Management Center Virtual is licensed appropriately.
- Perform these steps after the cluster is added to the Management Center Virtual:
  1. Configure platform settings with the health check port number in the Management Center. For more information on configuring this, see Platform Settings.
  2. Create a static route for data traffic. For more information on creating a static route, see Add a Static Route.

     Sample static route configuration:

```
Network: any-ipv4
Interface: vxlan_tunnel
Leaked from Virtual Router: Global
Gateway: vxlan_tunnel_gw
Tunneled: false
Metric: 2
```

Note *vxlan_tunnel_gw* is the data subnet's gateway IP address.

Deploy cluster on Azure with GWLB using an Azure resource manager template

This task creates a clustered deployment of Firewall Threat Defense Virtual instances on Azure using Gateway Load Balancer (GWLB) through Azure Resource Manager templates.

Deploy the Virtual Machine Scale Set for Azure GWLB using the customized Azure Resource Manager (ARM) template. The templates mentioned in the steps below are available on GitHub.

Procedure

Step 1 Prepare the template.

a) Clone the GitHub repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/azure.

b) Modify azure_ftdv_gwlb_cluster.json and azure_ftdv_gwlb_cluster_parameters.json with the required parameters.

Step 2 Log in to the Azure Portal: https://portal.azure.com.

Step 3 Create a Resource Group.

a) In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b) Choose the required **Region**.

Step 4 Create a virtual network with three subnets: Management, Data, and Cluster Control Link (CCL).

a) Create the virtual network.

1. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.
2. Choose the required **Region**. Click **Next: IP addresses**.

In the **IP Addresses** tab, click **Add subnet** and add the following subnets – Management, Data, and Cluster Control Link.

b) Add the subnets.

Step 5 Deploy the custom template.

a) Click **Create** > **Template deployment (deploy using custom templates)**.

b) Click **Build your own template in the editor**.

c) Click **Load File**, and upload **azure_ftdv_gwlb_cluster.json**.

d) Click **Save**.

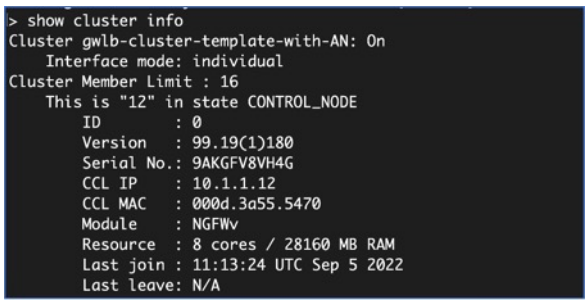
Step 6 Configure the Instance details.

a) Enter the required values and then click **Review + create**.

b) Click **Create** after the validation is passed.

c) After the instance is running, verify the cluster deployment by logging into any one of the nodes and entering the **show cluster info** command.

Figure 178: Show cluster info



Step 7 In the Azure Portal, click the Function app to register the cluster with the Cloud-Delivered Firewall Management Center.

Note
If you do not want to use the Function app, you can alternatively register the control node to the Cloud-Delivered Firewall Management Center directly by using **Add** > **Device** (not **Add** > **Cluster**). The rest of the cluster nodes will register automatically.

Step 8 Create FTPS Credentials by clicking **Deployment Center** > **FTPS credentials** > **User scope** > **Configure Username and Password**, and then click **Save**.

Step 9 Upload the Cluster_Function.zip file to the Function app by executing the following **curl** command in the local terminal.

curl -X POST -u *username* **--data-binary @"Cluster_Function.zip" https://** *Function_App_Name***.scm.azurewebsites.net/api/zipdeploy**

Note
The **curl** command might take few minutes (~2 to 3 minutes) to complete command execution.

The function will be uploaded to the Function app. The function will start, and you can see the logs in the storage account's outqueue. The device registration with the Management Center will be initiated.

Figure 179: Functions

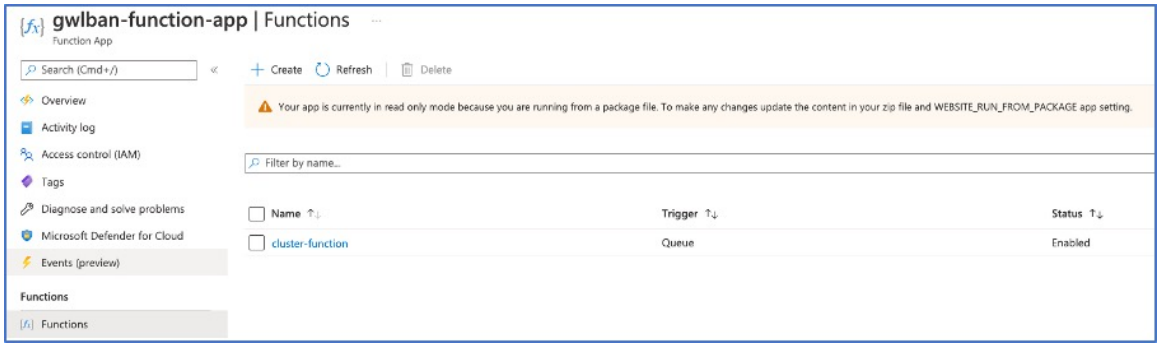


Figure 180: Queues

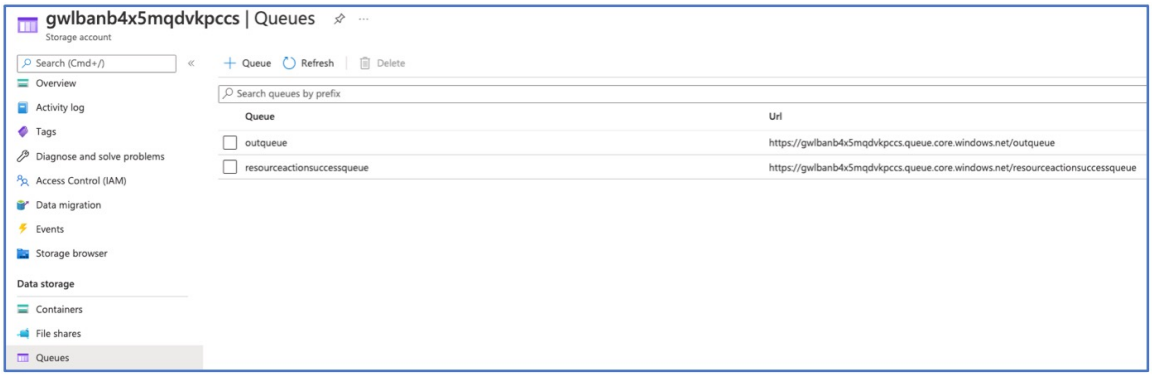
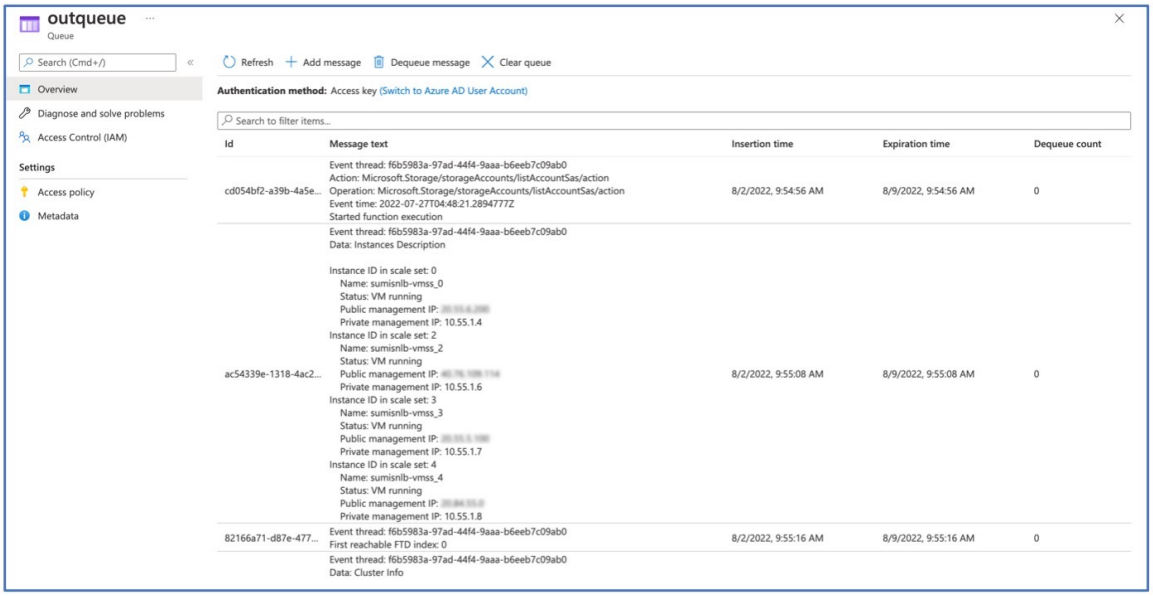


Figure 181: Outqueue



Sample topology for NLB-based cluster deployment

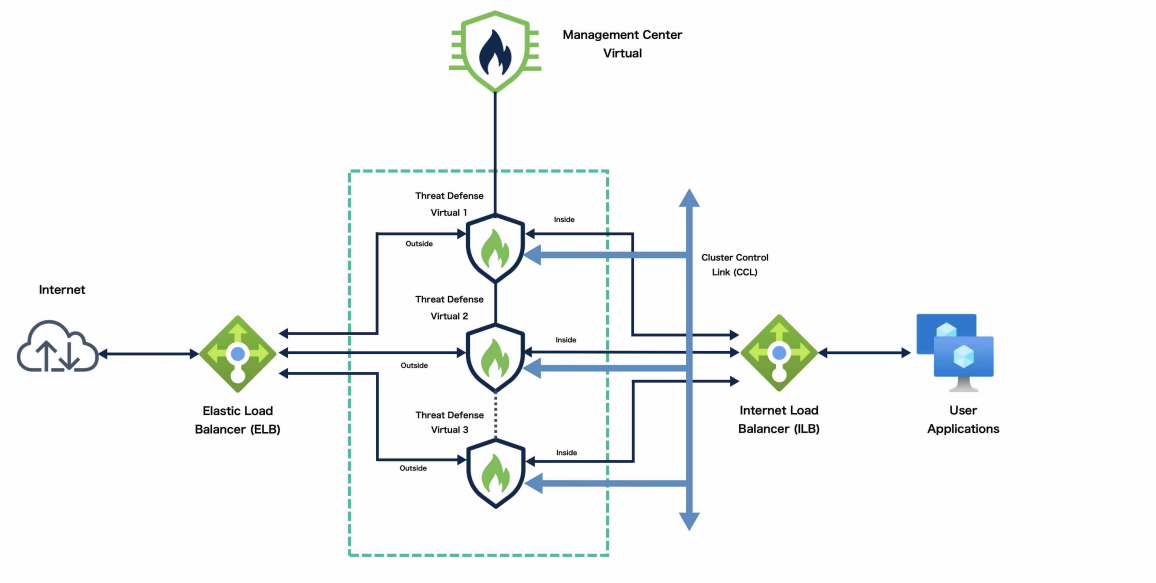
The sample topology provided here

- places the Threat Defense Virtual cluster between internal and external load balancers,
- uses a Management Center Virtual instance to manage the cluster, and
- depicts both inbound and outbound traffic flow patterns.

Traffic flow patterns

This topology shows how traffic moves through the network components:

- **Inbound traffic**: Traffic from the internet goes to the external load balancer, which transmits the traffic to the Threat Defense Virtual cluster. After the traffic has been inspected by a Threat Defense Virtual instance in the cluster, it is forwarded to the application virtual machine (VM).
- **Outbound traffic**: Traffic from the application VM is transmitted to the internal load balancer. Traffic is then forwarded to the Threat Defense Virtual cluster and then sent out to the internet.

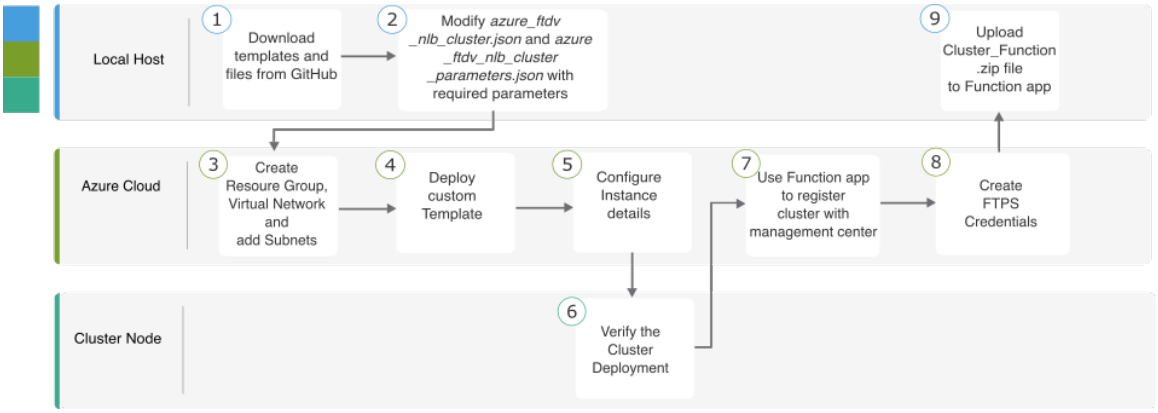


End-to-end process for deploying Threat Defense Virtual cluster in Azure with NLB

Provides a workflow to help users successfully deploy and verify a Threat Defense Virtual cluster on Azure with NLB using either template-based or manual deployment processes.

Template-based Deployment

The following flowchart illustrates the workflow of template-based deployment of Threat Defense Virtual cluster in Azure with NLB.

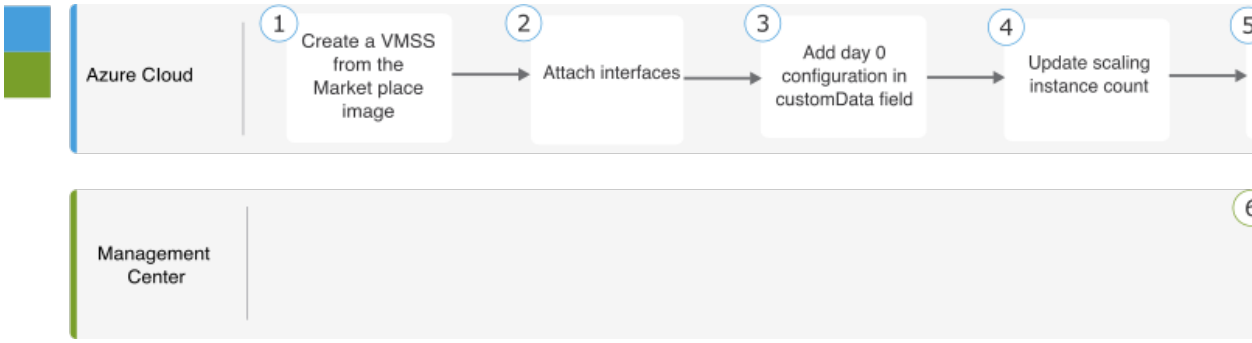


	Workspace	Steps
1	Local Host	Download templates and files from GitHub.

	Workspace	Steps
2	Local Host	Modify *azure_ftdv_nlb_cluster.json* and *azure_ftdv_nlb_cluster_parameters.json* with the required parameters.
3	Azure Cloud	Create the resource group, virtual network, and subnets.
4	Azure Cloud	Deploy custom template.
5	Azure Cloud	Configure instance details.
6	Cluster Node	Verify cluster deployment.
7	Azure Cloud	Use the Function app to register the cluster with the Management Center.
8	Azure Cloud	Create FTPS credentials.
9	Local Host	Upload *Cluster_Function.zip* file to the Function app.

Manual Deployment

The following flowchart illustrates the workflow of manual deployment of Threat Defense Virtual cluster in Azure with NLB.



	Workspace	Steps
1	Local Host	Create a VMSS from the Marketplace image.
2	Local Host	Attach interfaces.
3	Local Host	Add day 0 configuration in the customData field.
4	Local Host	Update scaling instance count.

	Workspace	Steps
5	Local Host	Configure NLB.
6	Management Center	Add control node.

Templates

The templates are available in GitHub. The parameter values are self-explanatory with the parameter names and values given in the template.

These templates provide the necessary components for cluster deployment:

Prerequisites for Azure NLB cluster deployment

Before deploying an Azure Network Load Balancer cluster, ensure that you complete these prerequisite tasks.

- To allow the cluster to auto-register with the Management Center, create a user with Network Admin & Maintenance User privileges on the Management Center. Users with these privileges can use REST API. Refer to the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the Management Center that matches the name of the policy that you will specify during template deployment.
- Ensure that the Management Center Virtual is licensed appropriately.
- After the cluster is added to the Management Center Virtual, complete these configuration steps:
  1. Configure platform settings with the health check port number in the Management Center. For more information on configuring this, refer to Platform Settings.
  2. Create static routes for traffic from outside and inside interfaces. For more information on creating a static route, refer to Add a Static Route.



Important Configure static routes for both interface gateway paths when Azure Load Balancer health probes are enabled on both inside and outside interfaces of FTDv cluster nodes. Do not rely on a single default route. Azure health probes may arrive on either interface, so FTDv must respond using the correct interface path. Missing routing configuration can cause probe drops and load balancer health check failures. If SSH on port 22 is used as the health probe and SSH is enabled on both interfaces through Platform Settings, verify return path routing for both interfaces. Similarly, if HTTP is used as the health probe and HTTP or HTTPS is enabled on both interfaces through Platform Settings, verify return path routing for both interfaces.

Note Ensure that you do not configure your virtual network using the Microsoft-owned IP address range that includes 168.63.129.16. This configuration causes a conflict with the health probe's IP address. For more information, refer to Microsoft Azure design guidance.

Sample static route configuration for the outside interface:

```
Network: any-ipv4
Interface: outside
Leaked from Virtual Router: Global
Gateway: ftdv-cluster-outside
Tunneled: false
Metric: 10
```

Note *ftdv-cluster-outside* is the outside subnet's gateway IP address.

Sample static route configuration for the inside interface:

```
Network: any-ipv4
Interface: inside
Leaked from Virtual Router: Global
Gateway: ftdv-cluster-inside-gw
Tunneled: false
Metric: 11
```

Note *ftdv-cluster-inside-gw* is the inside subnet's gateway IP address.

3. Configure NAT rule for data traffic. For more information on configuring NAT rules, refer to Network Address Translation.

Deploy cluster on Azure with NLB using an Azure resource manager template

This task enables you to deploy a cluster for Azure NLB using the customized Azure Resource Manager (ARM) template for scalable network security.

Deploy the cluster for Azure NLB using the customized Azure Resource Manager (ARM) template. Note that the templates mentioned in the steps below are available on GitHub.

Procedure

Step 1 Prepare the template.

a) Clone the GitHub repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/azure.

b) Modify azure_ftdv_nlb_cluster.json and azure_ftdv_nlb_cluster_parameters.json with the required parameters.

Step 2 Log in to the Azure Portal: https://portal.azure.com.

Step 3 Create a Resource Group.

a) In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.
b) Choose the required **Region**.

Step 4 Create a virtual network with 5 subnets: Management, Diagnostic, Inside, Outside, and Cluster Control Link.

a) Create the virtual network.

1. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.
2. Choose the required **Region**. Click **Next: IP addresses**.

b) Add the subnets.

In the **IP Addresses** tab, click **Add subnet** and add the following subnets – Management, Diagnostic, Inside, Outside, and Cluster Control Link.

Step 5 Deploy the custom template.

a) Click **Create** > **Template deployment (deploy using custom templates)**.
b) Click **Build your own template in the editor**.
c) Click **Load File**, and upload **azure_ftdv_nlb_cluster.json**.
d) Click **Save**.

Step 6 Configure the instance details.

a) Enter the required values and then click **Review + create**.

Note
For the cluster control link starting and ending addresses, specify only as many addresses as you need (up to 16). A larger range can affect performance.

b) Click **Create** after the validation is passed.
c) After the instance is running, verify the cluster deployment by logging into any one of the nodes and using the **show cluster info** command.

Figure 182: Show cluster info

```
> show cluster info
Cluster gwlb-cluster-template-with-AN: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "12" in state CONTROL_NODE
        ID          : 0
        Version     : 99.19(1)180
        Serial No.: 9AKGFV8VH4G
        CCL IP      : 10.1.1.12
        CCL MAC     : 000d.3a55.5470
        Module      : NGFWv
        Resource    : 8 cores / 28160 MB RAM
        Last join : 11:13:24 UTC Sep 5 2022
        Last leave: N/A
```

Step 7 In the Azure Portal, click the Function app to register the cluster to the Cloud-Delivered Firewall Management Center.

Note

If you do not want to use the Function app, you can alternatively register the control node with the Management Center directly by using **Add** > **Device** (not **Add** > **Cluster**). The rest of the cluster nodes will register automatically.

Step 8 Create FTPS Credentials by clicking **Deployment Center** > **FTPS credentials** > **User scope** > **Configure Username and Password**, and then click **Save**.

Step 9 Upload the Cluster_Function.zip file to the Function app by executing the **curl** command in the local terminal.

curl -X POST -u *username* **--data-binary @"Cluster_Function.zip" https://** *Function_App_Name***.scm.azurewebsites.net/api/zipdeploy**

Note
The **curl** command might take a few minutes (~2 to 3 minutes) to complete command execution.

The function will be uploaded to the Function app. The function will start, and you can see the logs in the storage account's outqueue. The device registration with the Management Center will be initiated.

Manual cluster deployment in Azure

To deploy the cluster manually, prepare the day0 configuration, deploy each node, and then add the control node to the Cloud-Delivered Firewall Management Center .

Create the Day0 configuration for Azure

You can use either a fixed configuration or a customized configuration.

Create the Day0 configuration with a fixed configuration for Azure

The fixed configuration will auto-generate the cluster bootstrap configuration.

```
"Cluster": {
"CclSubnetRange": "ip_address_start ip_address_end",
"ClusterGroupName": "cluster_name",
"HealthProbePort": "port_number",
"GatewayLoadBalancerIP": "ip_address",
"EncapsulationType": "vxlan",
"InternalPort": "internal_port_number",
"ExternalPort": "external_port_number",
"InternalSegId": "internal_segment_id",
"ExternalSegId": "external_segment_id"
}
```

Example

A sample day 0 configuration is given below.

```
"Cluster": {
"CclSubnetRange": "10.45.3.4 10.45.3.30",     //mandatory user input
"ClusterGroupName": "ngfwv-cluster",          //mandatory user input
"HealthProbePort": "7777",                    //mandatory user input
"GatewayLoadBalancerIP": "10.45.2.4",         //mandatory user input
"EncapsulationType": "vxlan",
"InternalPort": "2000",
```

```
"ExternalPort": "2001",
"InternalSegId": "800",
"ExternalSegId": "801"
}
```

Note If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration

For the Azure health check settings, be sure to specify the **HealthProbePort** you set here.

For the **CclSubnetRange** variable, specify a range of IP addresses starting from x.x.x.4. Ensure that you have at least 16 available IP addresses for clustering. Some examples of start and end IP addresses are given below.

Note All cluster infrastructure subnets must use /27 CIDR

Table 58: Examples of start and end IP addresses

CIDR	Start IP Address	End IP Address
10.1.1.0/27	10.1.1.4	10.1.1.30
10.1.1.32/27	10.1.1.36	10.1.1.62
10.1.1.64/27	10.1.1.68	10.1.1.94
10.1.1.96/27	10.1.1.100	10.1.1.126
10.1.1.128/27	10.1.1.132	10.1.1.158
10.1.1.160/27	10.1.1.164	10.1.1.190
10.1.1.192/27	10.1.1.196	10.1.1.222
10.1.1.224/27	10.1.1.228	10.1.1.254

Create the Day0 configuration with a customized configuration for azure

You can enter the entire cluster bootstrap configuration using commands.

```
"Cluster": {
"CclSubnetRange": "ip_address_start ip_address_end",
"ClusterGroupName": "cluster_name",
"HealthProbePort": "port_number",
"GatewayLoadBalancerIP": "ip_address",
"EncapsulationType": "vxlan",
"InternalPort": "internal_port_number",
"ExternalPort": "external_port_number",
"InternalSegId": "internal_segment_id",
"ExternalSegId": "external_segment_id"
}
```

Deploy cluster nodes manually - GWLB-based deployment

This task allows you to deploy cluster nodes so they form a cluster in Azure using Gateway Load Balancer-based deployment.

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Log in to the Azure Portal: https://portal.azure.com

Step 2 Create a Resource Group.

a. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b. Choose the required **Region**.

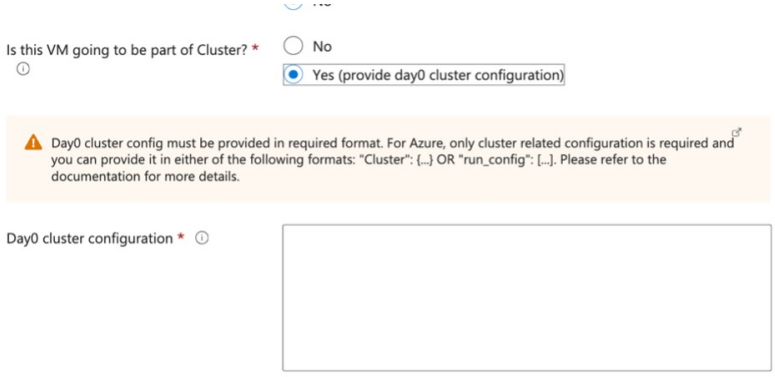
Step 3 Create a Virtual Network with the necessary subnets: Management, Data and Cluster Control Link (CCL).

Note
Configure the CCL with the smallest subnet mask as required. Wider subnets can impact performance.

Refer to the Azure document for creating the Virtual Network and subnet: https://learn.microsoft.com/en-us/azure/virtual-network/quickstart-create-virtual-network?tabs=portal

Step 4 Go to the Marketplace and search for **Cisco Secure Firewall Threat Defense Virtual – BYOL and PAYG** and click **Create**.

Step 5 Fill the required details and choose **Yes** for **Is this VM going to be part of Cluster?**



Paste the following cluster-related configuration in the text box.

```
"Cluster": {
"CclSubnetRange": "ip_address_start ip_address_end", //mandatory user input
"ClusterGroupName": "cluster_name", //mandatory user input
"HealthProbePort": "port_number", //mandatory user input
"GatewayLoadBalancerIP": "ip_address",  //mandatory user input
"EncapsulationType": "vxlan",
"InternalPort": "internal_port_number",
"ExternalPort": "external_port_number",
"InternalSegId": "internal_segment_id",
"ExternalSegId": "external_segment_id"
}
```

a) Click **Next** and select the **Virtual Network & Subnets**.
b) Click **Review + create**. Wait for the Threat Defense Virtual deployment to complete.

Step 6 Connect to the Threat Defense Virtual device and use the **show cluster info** command to confirm the cluster formation is successful.

```
> show cluster info
Cluster ngfwv-cluster: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "4" in state CONTROL_NODE
        ID        : 0
        Version   : 9.23(1)
        Serial No.: 9AC1VMGJKAQ
        CCL IP    : 169.254.200.4
        CCL MAC   : 6045.bda8.e07b
        Module    : NGFWv
        Resource  : 4 cores / 14336 MB RAM
        Last join : 05:22:55 UTC Jul 14 2025
        Last leave: N/A
Other members in the cluster:
    There is no other unit in the cluster
>
```

Step 7 Configure the Azure Gateway Load Balancer. Refer to Auto Scale with Azure Gateway Load Balancer Use Case for more information.

Step 8 Add the control node to the Cloud-Delivered Firewall Management Center. See Onboard a Secure Firewall Threat Defense Cluster.

Deploy cluster nodes manually - NLB-based deployment

Deploy cluster nodes manually to establish a functional cluster configuration in Azure using NLB-based deployment.

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Log in to the Azure Portal: https://portal.azure.com

Step 2 Create a Resource Group.

a. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b. Choose the required **Region**.

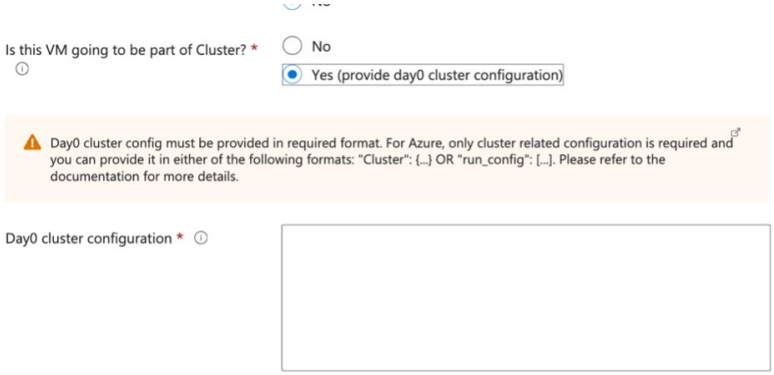
Step 3 Create a Virtual Network with the necessary subnets: Management, Inside, Outside and Cluster Control Link (CCL).

Note
Configure the CCL with the smallest subnet mask as required. Wider subnets can impact performance.

Refer to the Azure document for creating the Virtual Network and subnet: https://learn.microsoft.com/en-us/azure/virtual-network/quickstart-create-virtual-network?tabs=portal

Step 4 Go to the Marketplace and search for **Cisco Secure Firewall Threat Defense Virtual – BYOL and PAYG** and click **Create**.

Step 5 Fill the required details and choose **Yes** for **Is this VM going to be part of Cluster?**



Paste the following cluster-related configuration in the text box.

```
"Cluster": {
"CclSubnetRange": "ip_address_start ip_address_end", //mandatory user input
"ClusterGroupName": "cluster_name" //mandatory user input
}
```

Step 6 Click **Next** and select the **Virtual Network & Subnets**.

Step 7 Click **Review + create**. Wait until the Threat Defense Virtual deployment is completed.

Step 8 Connect to the Threat Defense Virtual device and use the **show cluster info** command to confirm the cluster formation is successful.

```
> show cluster info
Cluster ngfwv-cluster: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "4" in state CONTROL_NODE
        ID          : 0
        Version     : 9.23(1)
        Serial No.: 9AC1VMGJKAQ
        CCL IP      : 169.254.200.4
        CCL MAC     : 6045.bda8.e07b
        Module      : NGFWv
        Resource    : 4 cores / 14336 MB RAM
        Last join : 05:22:55 UTC Jul 14 2025
        Last leave: N/A
Other members in the cluster:
    There is no other unit in the cluster
>
```

Step 9 Add the control node to the Management Center. Refer to Onboard a Secure Firewall Threat Defense Cluster.

Troubleshoot cluster deployment in Azure

This reference provides troubleshooting guidance for resolving common issues that may occur during cluster deployment in Azure, including traffic flow problems, cluster formation failures, and role assignment errors.

- Issue: No traffic flow

Troubleshooting:

- Check if the health probe status of the Threat Defense Virtual instances deployed with a GWLB is healthy.
- If the Threat Defense Virtual instance's health probe status is unhealthy-
  - Check if the static route is configured in the Management Center Virtual.
  - Check if the default gateway is the data subnet's gateway IP address.
  - Check if the Threat Defense Virtual instance is receiving health probe traffic.
  - Check if the access list configured in the Management Center Virtual allows health probe traffic.

- Issue: Cluster is not formed

  Troubleshooting:

  - Check the IP address of the NVE-only cluster interface. Ensure that you can ping the NVE-only cluster interface of other nodes.
  - Check the IP address of the NVE-only cluster interfaces are part of the object group.
  - Ensure that the NVE interface is configured with the object group .
  - Ensure that the cluster interface in the cluster group has the right VNI interface. This VNI interface has the NVE with the corresponding object group.
  - Ensure that the nodes are pingable from each other. Since each node has its own cluster interface IP, these should be pingable from each other.
  - Check if the CCL subnet's start and end address mentioned during template deployment is correct. The start address should begin with the first available IP address in the subnet. For example, if the subnet is 192.168.1.0/27. The start address should be 192.168.1.4 (the three IP addresses at the start are reserved by Azure).
  - Check if the Management Center Virtual has a valid license.

- Issue: Role-related error while deploying resources again in the same resource group.

  Troubleshooting: Remove the roles given below by using these commands on the terminal.

  Error message:

```
"error": {
"code": "RoleAssignmentUpdateNotPermitted",
"message": "Tenant ID, application ID, principal ID, and scope are not allowed to be
updated."}
```

  - **az role assignment delete --resource-group** <*Resource Group Name*> **--role "Storage Queue Data Contributor"**
  - **az role assignment delete --resource-group** <*Resource Group Name*> **--role "Contributor"**

Deploy the cluster in GCP

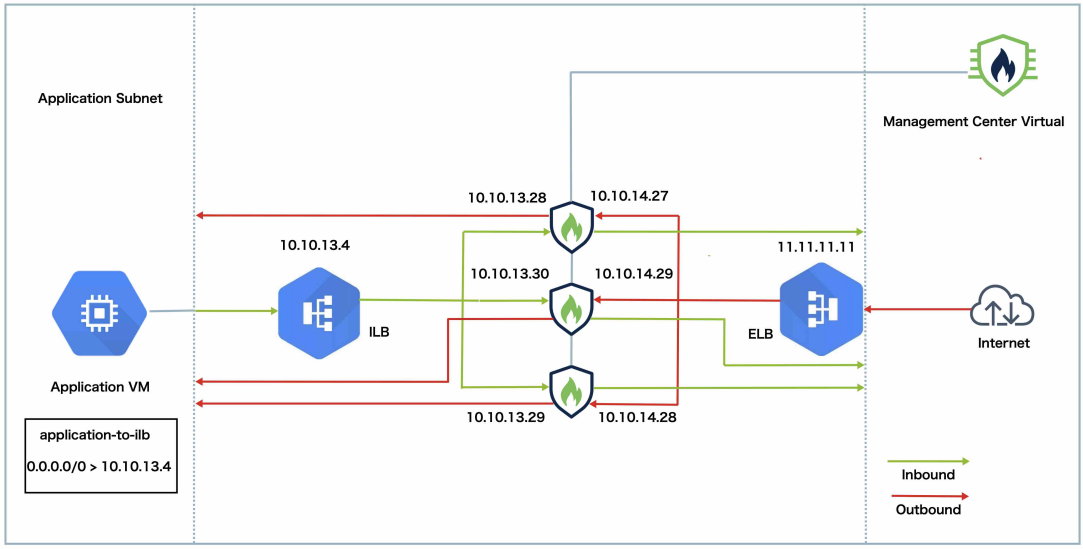
To deploy a cluster in GCP, you can either manually deploy or use an instance template to deploy an instance group. You can use the cluster with native GCP load balancers or non-native load balancers such as the Cisco Cloud Services Router.

Note Outbound traffic requires interface NAT and is limited to 64,000 connections.

Sample topology for GCP clustering autoscale solution

Understand the architecture and traffic flow of a typical GCP clustering autoscale deployment.

Figure 183: Sample Topology



The topology shows both inbound and outbound traffic flow.

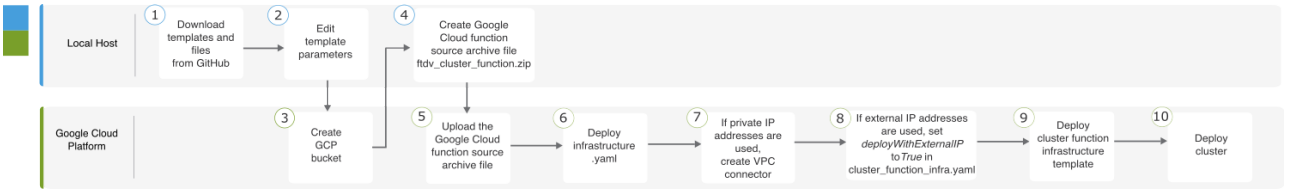
1. The Threat Defense Virtual cluster is placed between the internal and external load balancers. A Management Center Virtual instance is used to manage the cluster.
2. Inbound traffic from the internet goes to the external load balancer, which then transmits the traffic to the Threat Defense Virtual cluster.
3. The Threat Defense Virtual instance in the cluster inspects the traffic, and after inspection, forwards the traffic to the application VM.
4. Outbound traffic from the application VM goes to the internal load balancer. The load balancer forwards this traffic to the Threat Defense Virtual cluster, which sends it to the internet.

End-to-end process for deploying Threat Defense Virtual cluster in GCP

Provides a workflow to help users successfully deploy and verify a Threat Defense Virtual cluster on GCP using either template-based or manual deployment processes.

Template-based Deployment

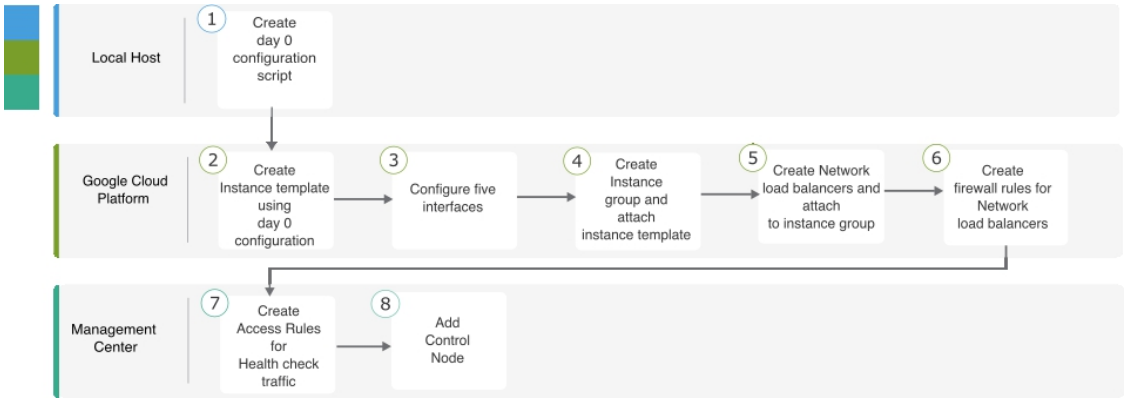
The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster on GCP.



	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Edit template parameters.
3	Google Cloud Platform	Create GCP bucket.
4	Local Host	Create Google Cloud function source archive file *ftdv_cluster_function.zip*.
5	Google Cloud Platform	Upload the Google function source archive file.
6	Google Cloud Platform	Deploy infrastructure.yaml.
7	Google Cloud Platform	If private IP addresses are used, create VPC connector.
8	Google Cloud Platform	If external IP addresses are used, set *deployWithExternalIP* to *True* in *cluster_function_infra.yaml*.
9	Google Cloud Platform	Deploy cluster function infrastructure template.
10	Google Cloud Platform	Deploy cluster.

Manual Deployment

The following flowchart illustrates the workflow for manual deployment of the Threat Defense Virtual cluster on GCP.



	Workspace	Steps
1	Local Host	Create day 0 configuration script.
2	Google Cloud Platform	Create instance template using day 0 configuration.
3	Google Cloud Platform	Configure the interfaces.
4	Google Cloud Platform	Create instance group and attach instance template.
5	Google Cloud Platform	Create NLB and attach to instance group.
6	Google Cloud Platform	Create firewall rules for NLB.
7	Management Center	Create access rules for health check traffic.
8	Management Center	Add control node.

Templates

The templates are available in GitHub. The parameter values are self-explanatory with the parameter names and values given in the template.

- Cluster deployment template for East-West traffic - deploy_ngfw_cluster_yaml
- Cluster deployment template for North-South traffic - deploy_ngfw_cluster.yaml

Deploy the instance group in GCP using an instance template

Deploy a scalable instance group in Google Cloud Platform using pre-configured instance templates to establish a Secure Firewall cluster with automated management capabilities.

Deploy the instance group in GCP using an instance template.

Before you begin

- Download the templates from GitHub to your local folder.
- Use Google Cloud Shell for deployment. Alternatively, you can use Google SDK on any macOS/Linux/Windows machine.
- To allow the cluster to auto-register with the Management Center, you need to create a user with administrative privileges on the Management Center that can use the REST API. See the *Cisco Secure Firewall Management Center Administration Guide*.
- Add an access policy in the Management Center that matches the name of the policy that you specified in *cluster_function_infra.yaml*.

Procedure

Step 1 Edit the **infrastructure.yaml**, **cluster_function_infra.yaml** and **deploy_ngfw_cluster.yaml** templates with the required *resourceNamePrefix* parameter (for example, ngfwvcls) and other required user inputs.

Note that there is a **deploy_ngfw_cluster.yaml** file in both the **east-west** and **north-south** folders in GitHub. Download the appropriate template as per your traffic flow requirement.

Step 2 Create a bucket using Google Cloud Shell to upload the Google cloud function source archive file *ftdv_cluster_function.zip*.

gsutil mb --pap enforced gs://*resourceNamePrefix***-ftdv-cluster-bucket/**

Ensure that the *resourceNamePrefix* variable here matches the *resourceNamePrefix* variable that you specified in **cluster_function_infra.yaml**.

Step 3 Create an archive file for the cluster infrastructure.

Example:

```
zip -j ftdv_cluster_function.zip ./cluster-function/*
```

Step 4 Upload the Google source archive that you created earlier.

gsutil cp ftdv_cluster_function.zip gs://*resourceNamePrefix***-ftdv-cluster-bucket/**

Step 5 Deploy infrastructure for the cluster.

gcloud deployment-manager deployments create *cluster_name* **--config infrastructure.yaml**

Step 6 If you are using private IP addresses, perform the steps given below:

a) Launch and set up the Management Center Virtual with a Threat Defense Virtual management VPC.

b) Create a VPC connector to connect the Google Cloud functions with the Threat Defense Virtual management VPC.

gcloud compute networks VPC-access connectors create *VPC-connector-name* **--region us-central1 --subnet** *resourceNamePrefix***-ftdv-mgmt-subnet28**

Step 7 If the Management Center is remote from the Threat Defense Virtual, and the Threat Defense Virtual needs an external IP address, ensure that you set **deployWithExternalIP** to **True** in **cluster_function_infra.yaml**.

Step 8 Deploy the cluster function infrastructure.

gcloud deployment-manager deployments create *cluster_name* **--config cluster_function_infra.yaml**

Step 9 Deploy the cluster.

a. For North-South topology deployment:

gcloud deployment-manager deployments create *cluster_name* **--config north-south/deploy_ngfw_cluster.yaml**

b. For East-West topology deployment:

gcloud deployment-manager deployments create *cluster_name* **--config east-west/deploy_ngfw_cluster.yaml**

Manual cluster deployment in GCP

To deploy the cluster manually, prepare the day0 configuration, deploy each node, and then add the control node to the Cloud-Delivered Firewall Management Center .

Create the Day0 configuration with a fixed configuration for GCP

The fixed configuration will auto-generate the cluster bootstrap configuration.

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
                "Diagnostic": "OFF",
    "Cluster": {
        "CclSubnetRange": "ip_address_start ip_address_end",
        "ClusterGroupName": "cluster_name"
    }
}
```

For example:

```
{
    "AdminPassword": "DeanW1nche$ter",
    "Hostname": "ciscoftdv",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "Cluster": {
        "CclSubnetRange": "10.10.55.2 10.10.55.253",     //mandatory user input
        "ClusterGroupName": "ftdv-cluster"               //mandatory user input
    }
}
```

Note If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration.

For the **CclSubnetRange** variable, note that you cannot use the first two IP addresses and the last two IP addresses in the subnet. See Reserved IP addresses in IPv4 subnets for more information. Ensure that you have at least 16 available IP addresses for clustering. Some examples of start and end IP addresses are given below.

Note All cluster infrastructure subnets must use /27 CIDR.

Table 59: Examples of start and end IP addresses

CIDR	Start IP Address	End IP Address
10.1.1.0/27	10.1.1.2	10.1.1.29
10.1.1.32/27	10.1.1.34	10.1.1.61
10.1.1.64/27	10.1.1.66	10.1.1.93
10.1.1.96/27	10.1.1.98	10.1.1.125
10.1.1.128/27	10.1.1.130	10.1.1.157
10.1.1.160/27	10.1.1.162	10.1.1.189
10.1.1.192/27	10.1.1.194	10.1.1.221
10.1.1.224/27	10.1.1.226	10.1.1.253

Deploy cluster nodes manually

Deploy cluster nodes to form a functional cluster on Google Cloud Platform (GCP) with proper network load balancing and management integration.

Deploy the cluster nodes so they form a cluster. For clustering on GCP, you cannot use the 4 vCPU machine type. The 4 vCPU machine type only supports four interfaces, and five are needed. Use a machine type that supports five interfaces, such as c2-standard-8.

Procedure

Step 1 Create an instance template using the day 0 configuration (in the **Metadata** > **Startup Script** section) with 5 interfaces: outside, inside, management, diagnostic, and cluster control link.

See Secure Firewall Threat Defense Virtual getting started guides.

Step 2 Create an instance group, and attach the instance template.

Step 3 Create GCP network load balancers (internal and external), and attach the instance group.

Step 4 For GCP network load balancers, allow health checks in your security policy on the Management Center. Refer to Allow health checks for GCP network load balancers, on page 897.

Step 5 Add the control node to the Management Center. Refer to Onboard a Secure Firewall Threat Defense Cluster.

Allow health checks for GCP network load balancers

Google Cloud provides health checks to determine if backends respond to traffic. This configuration allows health check traffic through firewall rules and access rules, redirects health check traffic to the Google metadata server using dynamic manual NAT rules, and sets up routing for GCP health checks across all interfaces used for health probes.

Configuration requirements

Refer to https://cloud.google.com/load-balancing/docs/health-checks to create firewall rules for network load balancers. Then in the Cloud-Delivered Firewall Management Center, create access rules to allow the health check traffic. See https://cloud.google.com/load-balancing/docs/health-check-concepts for the required network ranges. See Access Control Rules, on page 2041.

You also need to configure dynamic manual NAT rules to redirect the health check traffic to the Google metadata server at 169.254.169.254. See Configure dynamic manual NAT, on page 1596.

You can set up a route for GCP health checks across all interfaces that are used to configure their health probes. You can achieve this by creating a route with a higher metric on interfaces where a route for GCP health checks is not already available.

North-South NAT rules sample configuration

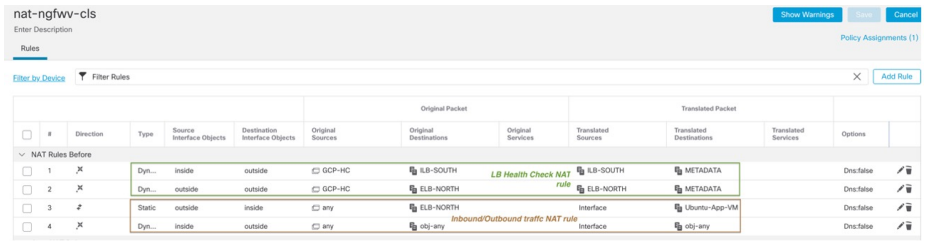
```
nat (inside,outside) source dynamic GCP-HC ILB-SOUTH destination static ILB-SOUTH METADATA
nat (outside,outside) source dynamic GCP-HC ELB-NORTH destination static ELB-NORTH METADATA

nat (outside,inside) source static any interface destination static ELB-NORTH Ubuntu-App-VM
nat (inside,outside) source dynamic any interface destination static obj-any obj-any

object network Metadata
 host 169.254.169.254

object network ILB-SOUTH
host <ILB_IP>
object network ELB-NORTH
host <ELB_IP>

object-group network GCP-HC
 network-object 35.191.0.0 255.255.0.0
 network-object 130.211.0.0 255.255.252.0
 network-object 209.85.204.0 255.255.252.0
 network-object 209.85.152.0 255.255.252.0
```



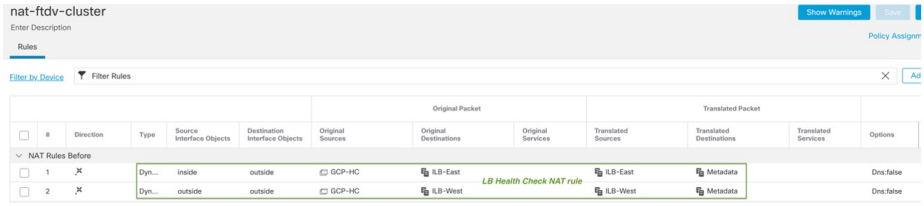
East-West NAT rules sample configuration

```
nat (inside,outside) source dynamic GCP-HC ILB-East destination static ILB-East Metadata
nat (outside,outside) source dynamic GCP-HC ILB-West destination static ILB-West Metadata
```

```
object network Metadata
 host 169.254.169.254

object network ILB-East
host <ILB_East_IP>
object network ILB-West
host <ILB_West_IP>

object-group network GCP-HC
 network-object 35.191.0.0 255.255.0.0
 network-object 130.211.0.0 255.255.252.0
 network-object 209.85.204.0 255.255.252.0
 network-object 209.85.152.0 255.255.252.0
```



North-South and East-West traffic routing configuration sample

```
route outside 0.0.0.0 0.0.0.0 <Outside_Gateway> 1
route inside 35.191.0.0 255.255.0.0 <Inside_Gateway> 1
route inside 130.211.0.0 255.255.252.0 <Inside_Gateway> 1
route inside 209.85.152.0 255.255.252.0 <Inside_Gateway> 1
route inside 209.85.204.0 255.255.252.0 <Inside_Gateway> 1
```

If a default route is not available, then policy-based routing can be used to route the traffic for health checks.

Note Set the NAT > Translated destination port to 80.

Add the cluster to the Management Center (manual deployment)

This task registers a manually deployed cluster with the Cloud-Delivered Firewall Management Center to enable centralized management and configuration of the cluster units.

Use this procedure to add the cluster to the Cloud-Delivered Firewall Management Center if you manually deployed the cluster. If you used a template, the cluster will auto-register on the Cloud-Delivered Firewall Management Center.

Add one of the cluster units as a new device to the Cloud-Delivered Firewall Management Center; the Cloud-Delivered Firewall Management Center automatically detects all other cluster members.

Before you begin

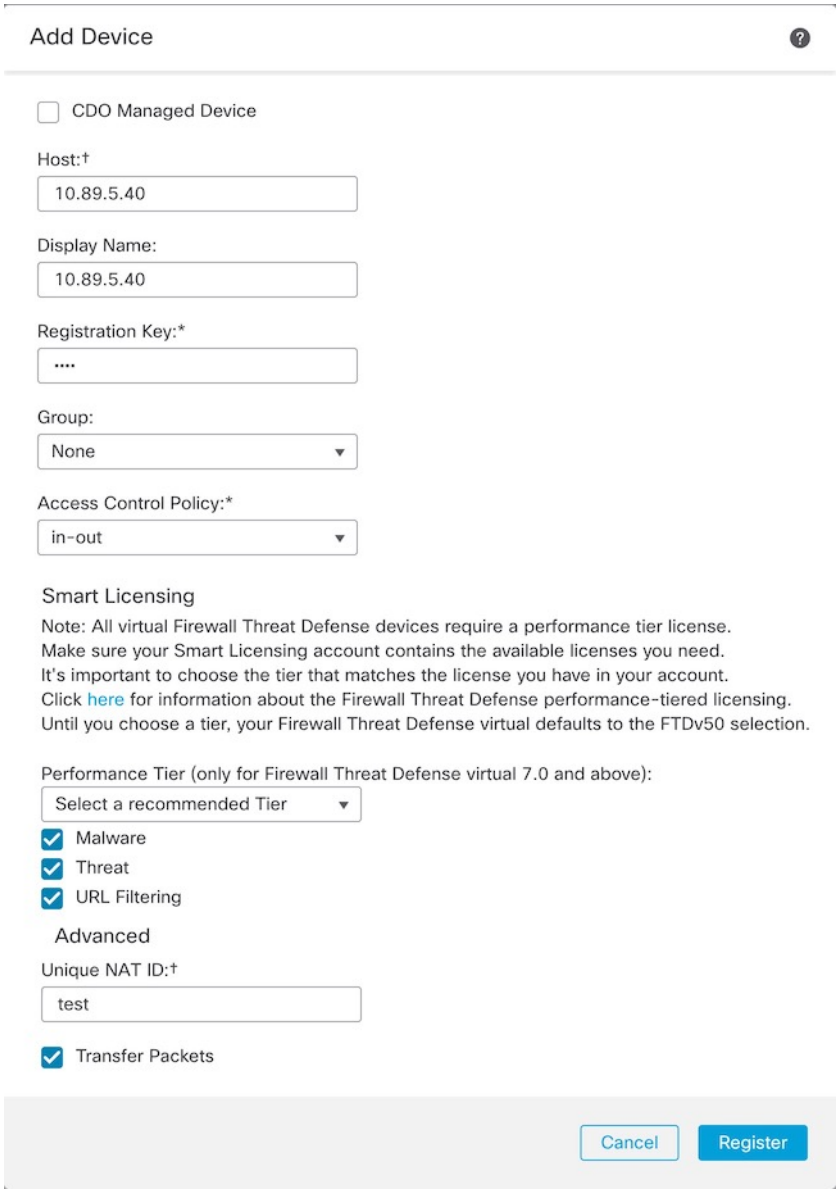
- All cluster units must be in a successfully formed cluster before adding the cluster to the Cloud-Delivered Firewall Management Center. You should also check which unit is the control unit. Use the Firewall Threat Defense **show cluster info** command.

Follow these steps to add the cluster to the Management Center for manual deployment:

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and then choose **Add** > **Add Device** to add the control unit using the unit's management IP address.

Figure 184: Add device



a) In the **Host** field, enter the IP address or hostname of the control unit.

We recommend adding the control unit for the best performance, but you can add any unit of the cluster.

If you used a NAT ID during device setup, you may not need to enter this field.

b) In the **Display Name** field, enter a name for the control unit as you want it to display in the Cloud-Delivered Firewall Management Center.

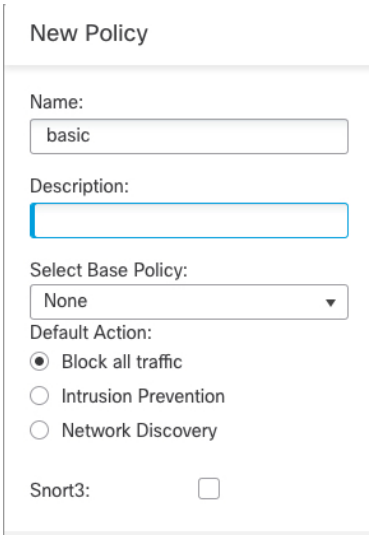
This display name is not for the cluster; it is only for the control unit you are adding. You can later change the name of other cluster members and the cluster display name.

c) In the **Registration Key** field, enter the same registration key that you used during device setup. The registration key is a one-time-use shared secret.

d) (Optional) Add the device to a device **Group**.

e) Choose an initial **Access Control Policy** to deploy to the device upon registration, or create a new policy.

If you create a new policy, you create a basic policy only. You can later customize the policy as needed.



f) Choose licenses to apply to the device.

g) If you used a NAT ID during device setup, expand the **Advanced** section and enter the same NAT ID in the **Unique NAT ID** field.

h) Check the **Transfer Packets** check box to allow the device to transfer packets to the Cloud-Delivered Firewall Management Center.

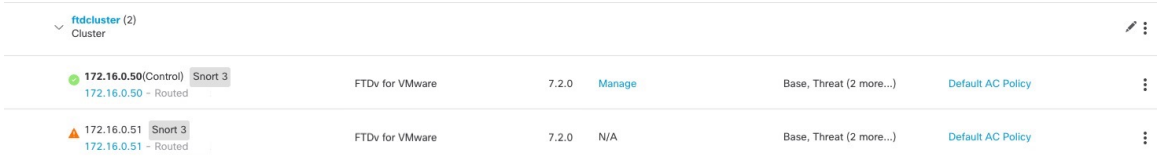
This option is enabled by default. When events such as IPS or Snort are triggered with this option enabled, the device sends event metadata information and packet data to the Cloud-Delivered Firewall Management Center for inspection. If you disable it, only event information will be sent to the Cloud-Delivered Firewall Management Center, but packet data is not sent.

i) Click **Register**.

The Cloud-Delivered Firewall Management Center identifies and registers the control unit, and then registers all data units. If the control unit does not successfully register, then the cluster is not added. A registration failure can occur if the cluster was not up, or because of other connectivity issues. In this case, we recommend that you try re-adding the cluster unit.

The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster units.

Figure 185: Cluster management

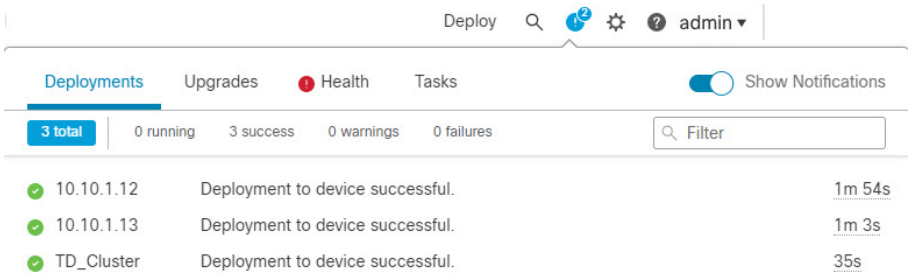


A unit that is currently registering shows the loading icon.

Figure 186: Node registration



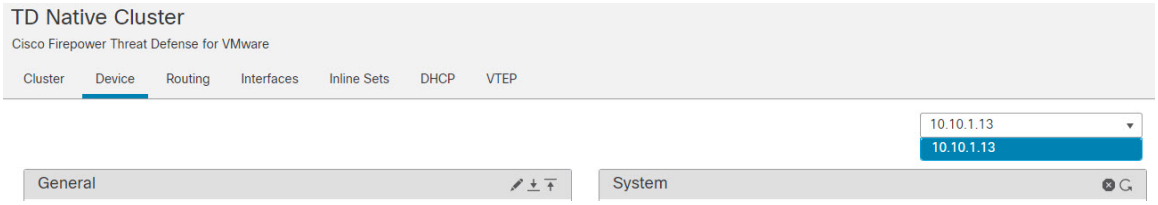
You can monitor cluster unit registration by clicking the **Notifications** icon and choosing **Tasks**. The Cloud-Delivered Firewall Management Center updates the Cluster Registration task as each unit registers. If any units fail to register, see Reconcile Cluster Nodes, on page 911.



Step 2 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

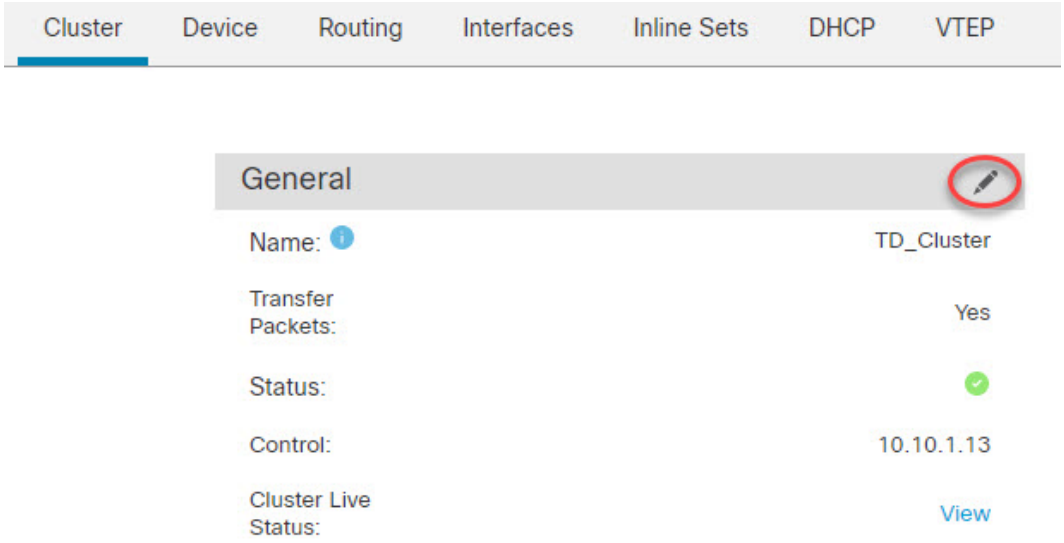
Most configuration can be applied to the cluster as a whole, and not nodes in the cluster. For example, you can change the display name per node, but you can only configure interfaces for the whole cluster.

Step 3 On the **Devices** > **Device Management** and then choose **Add**, **Cluster** screen, you see **General**, **License**, **System**, and **Health** settings.

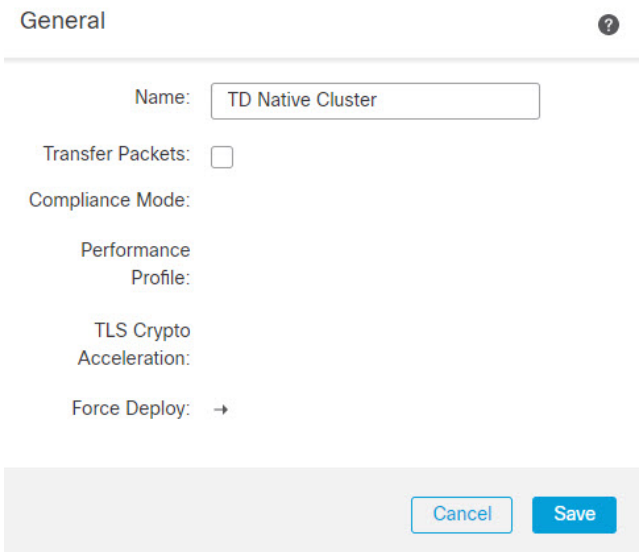


See these cluster-specific items:

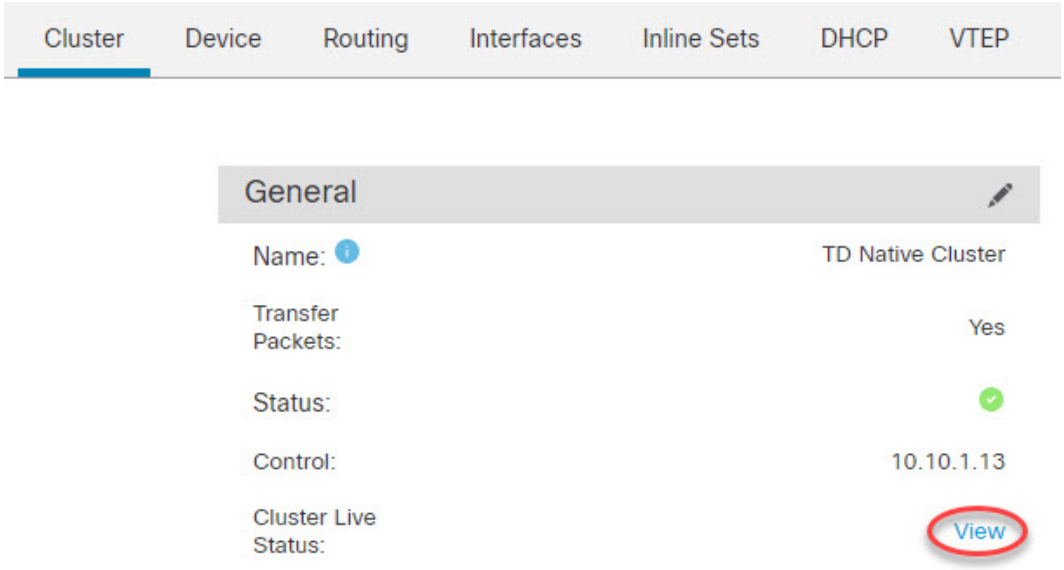
- **General > Name**—Change the cluster display name by clicking the **Edit** (✎).



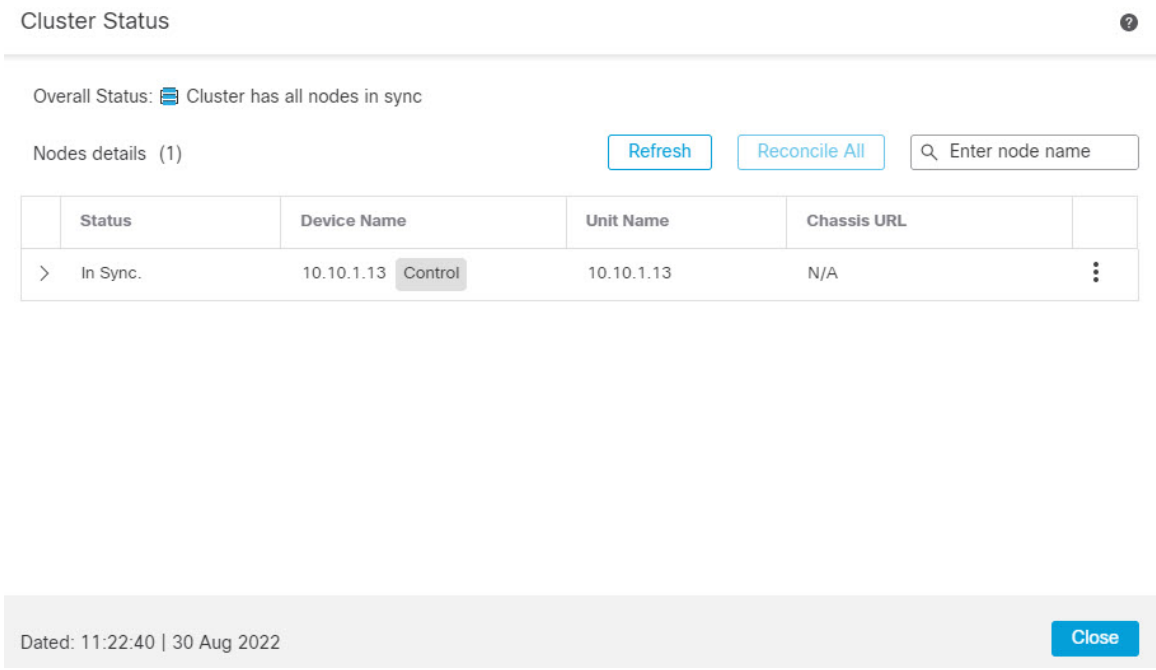
Then set the **Name** field.



- **General > Cluster Live Status**—Click the **View** link to open the **Cluster Status** dialog box.



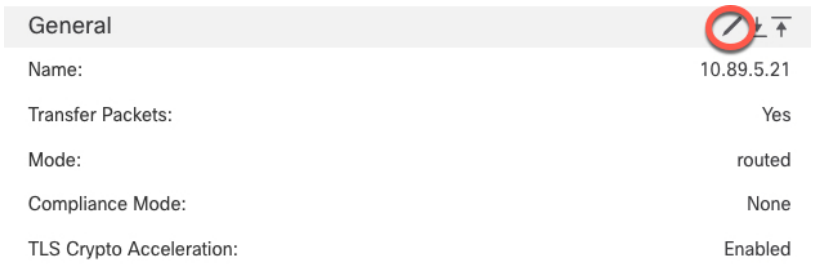
The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile**.



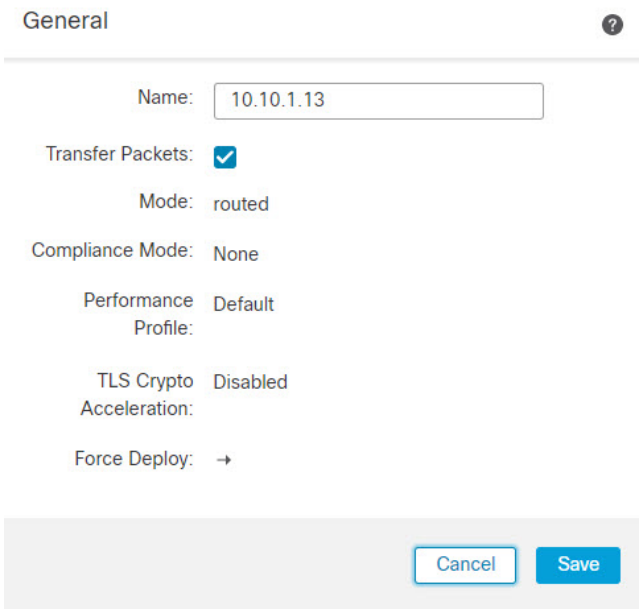
- **License**—Click **Edit** (✎) to set license entitlements.

Step 4 On the **Devices** > **Device Management** and then click **Add** > **Device**, you can choose each member in the cluster from the top right drop-down menu and configure these settings.

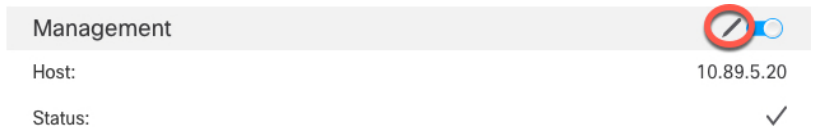
- **General > Name**—Change the cluster member display name by clicking the **Edit** (✎).



Then set the **Name** field.



- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the Cloud-Delivered Firewall Management Center so that it can reach the device on the network; edit the **Host** address in the **Management** area.



Onboard a Secure Firewall Threat Defense Cluster

This task enables you to register a Secure Firewall Threat Defense cluster with Cloud-Delivered Firewall Management Center for centralized management and policy enforcement.

The following table provides information about device models that support cluster onboarding and creation on the Cloud-Delivered Firewall Management Center:

Secure Firewall Threat Defense Platforms	Minimum Secure Firewall Threat Defense Version for Cluster Management	Support cluster creation from Cloud-Delivered Firewall Management Center?
VMware, KVM	7.2.1	Yes
AWS, GCP	7.2.1	No
Azure	7.3	No
Secure Firewall 3100	7.2.1	Yes
Firepower 4100	7.2.x	No
Secure Firewall 4200	7.4	Yes
Firepower 9300	7.2.x	No

Note When you onboard a Secure Firewall Threat Defense cluster to Cloud-Delivered Firewall Management Center, you register only the cluster CONTROL NODE. After the CONTROL NODE registers, Cloud-Delivered Firewall Management Center discovers the DATA nodes in the cluster.

Before you begin

Review these limitations:

- Firepower 4100 and Firepower 9300 devices must be clustered through the device's Firewall Chassis Manager.
- Secure Firewall 3100 devices, Secure Firewall 4200 devices, KVM, and VMware environments must be clustered through the Cloud-Delivered Firewall Management Center UI.
- Azure, AWS, and GCP environment clusters must be created through their own environment and onboarded to Cloud-Delivered Firewall Management Center.
- Verify that the Secure Firewall Threat Defense virtual devices are successfully configured as a cluster before you register the cluster with Cloud-Delivered Firewall Management Center.

  For example, in the Secure Firewall Threat Defense CLI, enter the `show cluster info` command. Confirm that the output shows the cluster as `On`, one unit in `CONTROL_NODE` state, and the remaining units in `DATA_NODE` state.

```
> show cluster info
Cluster mAZ-ngfw-cl: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "74-a" in state DATA_NODE
        ID        : 2
        Version   : 9.22(1)1
        Serial No.: 9AUVQ3DSF66
        CCL IP    : 1.1.1.74
        CCL MAC   : 02e2.778f.d3ed
        Module    : NGFWv
        Resource  : 4 cores / 7680 MB RAM
        Last join : 07:28:26 UTC Sep 25 2024
        Last leave: 07:28:11 UTC Sep 25 2024
Other members in the cluster:
    Unit "135-b" in state CONTROL_NODE
        ID        : 0
        Version   : 9.22(1)1
        Serial No.: 9A6W0A51KGK
        CCL IP    : 1.1.2.135
        CCL MAC   : 1294.34ae.4ce9
        Module    : NGFWv
        Resource  : 4 cores / 7680 MB RAM
        Last join : 09:45:52 UTC Sep 24 2024
        Last leave: N/A
    Unit "183-c" in state DATA_NODE
        ID        : 1
        Version   : 9.22(1)1
        Serial No.: 9A1S400HL8F
        CCL IP    : 1.1.3.183
        CCL MAC   : 0aff.e889.f193
        Module    : NGFWv
        Resource  : 4 cores / 7680 MB RAM
        Last join : 07:29:29 UTC Sep 25 2024
        Last leave: 07:28:11 UTC Sep 25 2024
>
```

Follow these steps to onboard a Secure Firewall Threat Defense cluster:

Procedure

Step 1 In the left pane, click **Security Devices**.

Step 2 Click **Onboard device or service** ( ) icon.

Step 3 Click the **FTD** tile.

Step 4 Under **Management Mode**, select **FTD**. By selecting **FTD** under **Management Mode**, you will not be able to manage the device using the previous management platform. All existing policy configurations except for interface configurations will be reset. You must re-configure policies after you onboard the device.

Step 5 Select **Use CLI Registration Key**.

Step 6 Enter the device name in the **Device Name** field and click **Next**.

Step 7 In the **Policy Assignment** step, use the drop-down menu to select an access control policy to deploy once the device is onboarded. If you have no policies configured, select the **Default Access Control Policy**.

Step 8 Specify whether the device you are onboarding is a physical or virtual device. If you are onboarding a virtual device, you must select the device's performance tier from the drop-down menu.

Step 9 Select the subscription licenses you want to apply to the device. Click **Next**.

Step 10 Security Cloud Control generates a command with the registration key. Paste the entire registration key into the device's CLI without making changes.

Step 11 (Optional) Add labels to your device to help sort and filter the **Security Devices** page. Enter a label and select the blue plus button. Labels are applied to the device after it's onboarded to Security Cloud Control.

Once the device is synchronized, select the device you just onboarded from the **Security Devices** page and select any of the options listed under the **Device Management** pane located to the right. We strongly recommend the following actions:

- If you did not already, create a custom access control policy to customize the security for your environment. See Access Control Overview in *Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control* for more information.
- Enable Cisco Security Analytics and Logging (SAL) to view events in the Security Cloud Control dashboard **or** register the device to an Secure Firewall Management Center for security analytics. See Cisco Security Analytics and Logging in *Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control* for more information.

Edit cluster health monitor settings

This task allows you to modify cluster health monitor settings to control how cluster nodes monitor system health and automatically rejoin the cluster after failures.

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings for monitoring cluster node health. You can monitor any port-channel ID, any single physical interface ID, and the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link. It is always monitored.

When a topology change occurs (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature. Also, disable interface monitoring for any interfaces that are disabled. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Procedure

Step 1 Choose **Devices** > **Device Management**.

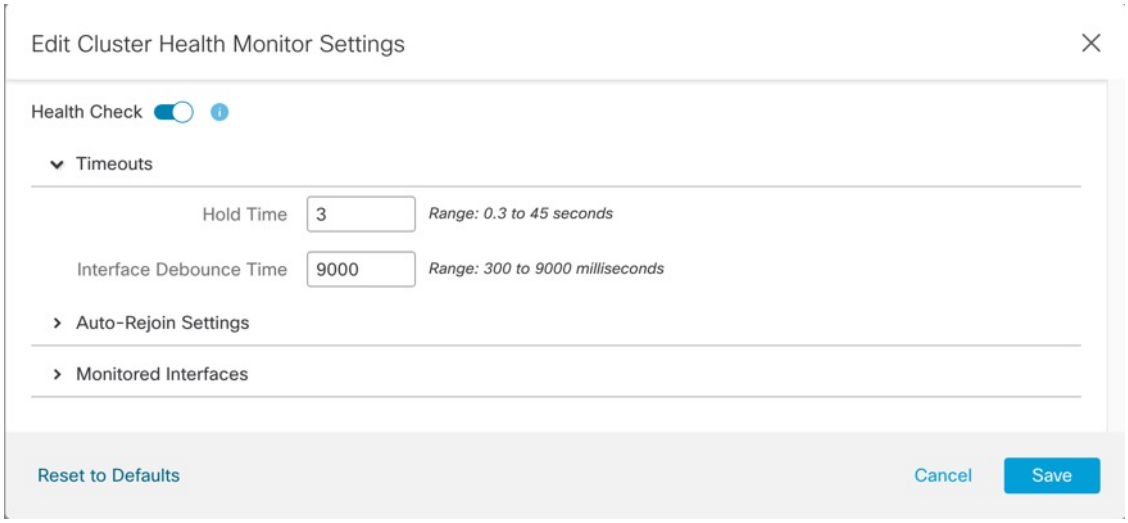
Step 2 Next to the cluster you want to modify, click **Edit** ( ).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** ( ).

Step 5 Disable the system health check by clicking the **Health Check** slider.

Figure 187: Disable the system health check



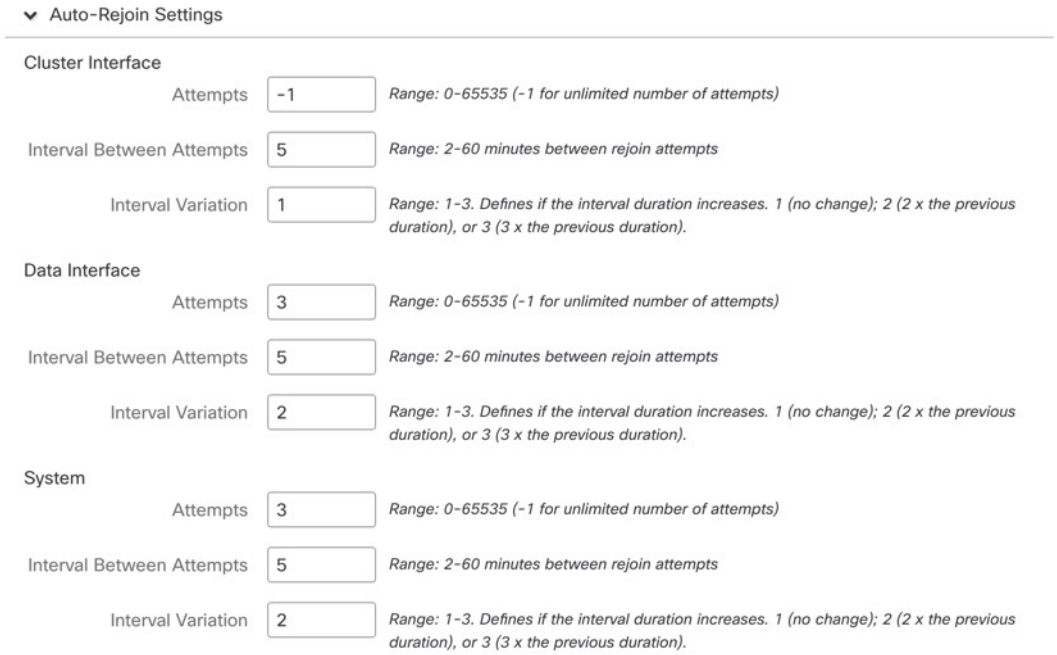
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between 0.3 and 45 seconds. The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the specified number of milliseconds before marking the interface as failed and removing the node from the cluster. If an EtherChannel transitions from a down state to an up state, such as after the switch reloads or enables an EtherChannel, a longer debounce time can prevent the interface from appearing failed on a cluster node if another cluster node bundles the ports more quickly.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 188: Configure Auto-Rejoin settings

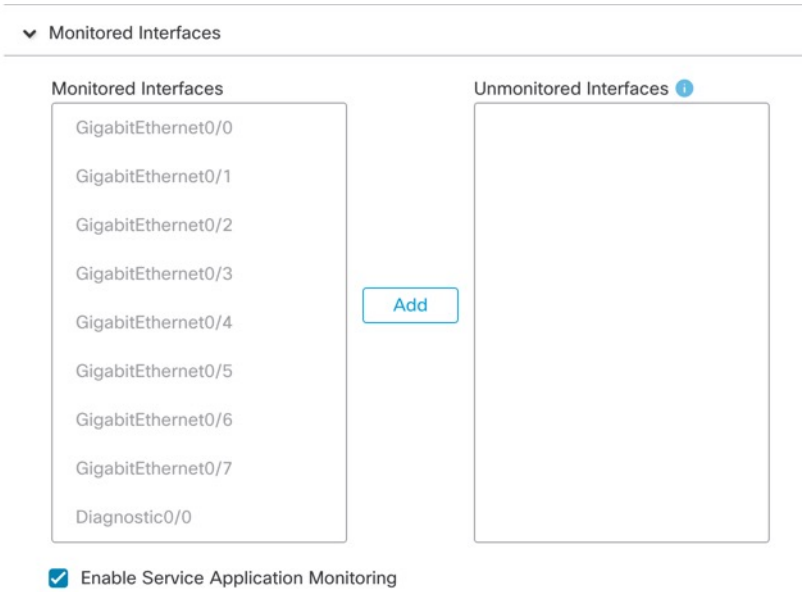


Set values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 189: Configure monitored interfaces



The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The time required to remove a member from the cluster depends on the type of interface and whether the node is established or is joining the cluster. By default, health check is enabled for all interfaces and the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Deploy configuration changes.

Manage cluster nodes

Explains how to add, remove, and configure cluster nodes in your deployment.

Disable clustering

Disable clustering to temporarily deactivate a node in preparation for deleting the node or for maintenance purposes.

You may want to deactivate a node to delete it or for temporary maintenance. This procedure temporarily deactivates a node, but the node will still appear in the Cloud-Delivered Firewall Management Center device list. When a node becomes inactive, all data interfaces are shut down.

Note Do not power off the node without first disabling clustering.

Procedure

Step 1 For the unit you want to disable, choose **Devices** > **Device Management**, click the **More** (⋮), and choose **Disable Node Clustering**.

Step 2 Confirm that you want to disable clustering on the node.

The node will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.

Step 3 To reenable clustering, refer to Rejoin the Cluster.

Rejoin the cluster

Rejoin a node to the cluster after it was removed because of failure or when clustering was manually disabled.

If a node was removed from the cluster, for example, due to a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management**, click **More** (⋮), and choose **Enable Node Clustering**.

Step 2 Confirm that you want to enable clustering on the node.

Reconcile Cluster Nodes

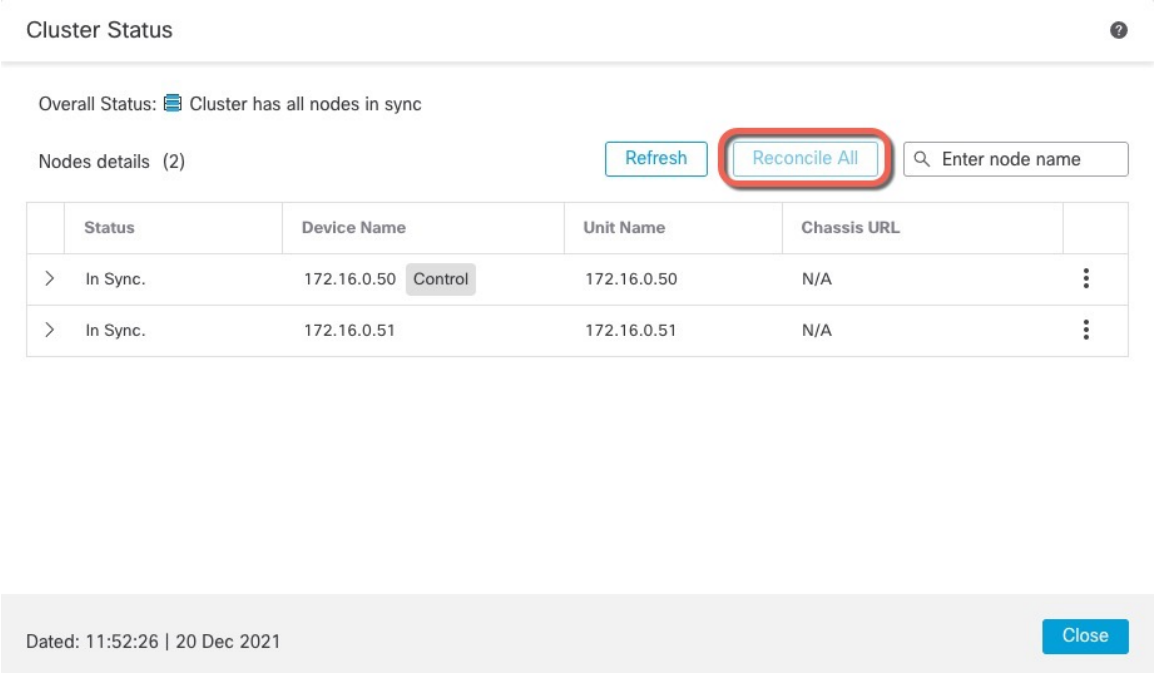
If a cluster node fails to register, you can reconcile the cluster membership from the device to the Cloud-Delivered Firewall Management Center. For example, a data node might fail to register if the Cloud-Delivered Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

Step 1 Choose **Devices** > **Device Management More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Step 2 Click **Reconcile All**.

Figure 190: Reconcile All



Unregister the cluster or nodes and register with a new Cloud-Delivered Firewall Management Center

This task allows you to move cluster nodes from one management center to another by safely removing them from their current management center and re-registering them with a new one.

You can unregister the cluster from the Cloud-Delivered Firewall Management Center, which keeps the cluster intact. You might want to unregister the cluster if you want to add the cluster to a new Cloud-Delivered Firewall Management Center.

You can also unregister a node from the Cloud-Delivered Firewall Management Center without breaking the node from the cluster. Although the node is not visible in the Cloud-Delivered Firewall Management Center, it is still part of the cluster, and it will continue to pass traffic and could even become the control node. You cannot unregister the current control node. You might want to unregister the node if it is no longer reachable from the Cloud-Delivered Firewall Management Center, but you still want to keep it as part of the cluster while you troubleshoot management connectivity.

Unregistering a cluster:

- Severs all communication between the Cloud-Delivered Firewall Management Center and the cluster.
- Removes the cluster from the **Device Management** page.
- Returns the cluster to local time management if the cluster's platform settings policy is configured to receive time from the Cloud-Delivered Firewall Management Center using NTP.
- Leaves the configuration intact, so the cluster continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the cluster again to the same or a different Cloud-Delivered Firewall Management Center causes the configuration to be removed, so the cluster will stop processing traffic at that point; the cluster configuration remains intact so you can add the cluster as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.

Before you begin

This procedure requires CLI access to one of the nodes.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster or node, and choose **Unregister**.

Step 2 You are prompted to unregister the cluster or node; click **Yes**.

Step 3 You can register the cluster with a new (or the same) Cloud-Delivered Firewall Management Center by adding one of the cluster members as a new device.

Add one of the cluster nodes as a device. The system will discover the remaining nodes.

Step 4 To re-add a deleted node, refer to Reconcile Cluster Nodes, on page 911.

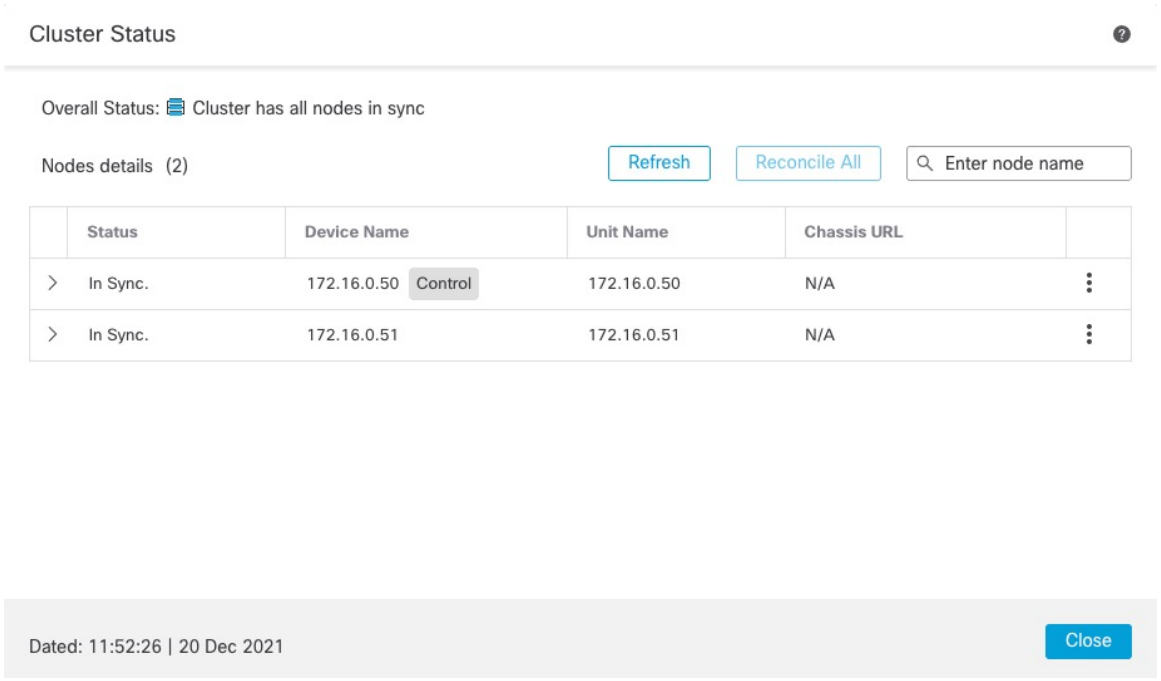
Monitor cluster status

You can monitor the cluster in the Cloud-Delivered Firewall Management Center and at the Firewall Threat Defense CLI.

Monitor cluster status through these methods:

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management**, **More** (⋮) icon or from the **Devices** > **Device Management**, click **Add**, choose the **Cluster** page **General** area **Cluster Live Status** link.

Figure 191: Cluster status



The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes these states:

- In Sync—The node is registered with the Cloud-Delivered Firewall Management Center.
- Pending Registration—The node is part of the cluster, but has not yet registered with the Cloud-Delivered Firewall Management Center. If a node fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The node is registered with the Cloud-Delivered Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the node from the cluster.
- Joining cluster...—The node is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Cloud-Delivered Firewall Management Center.

For each node, you can view the **Summary** or the **History**.

Figure 192: Node summary

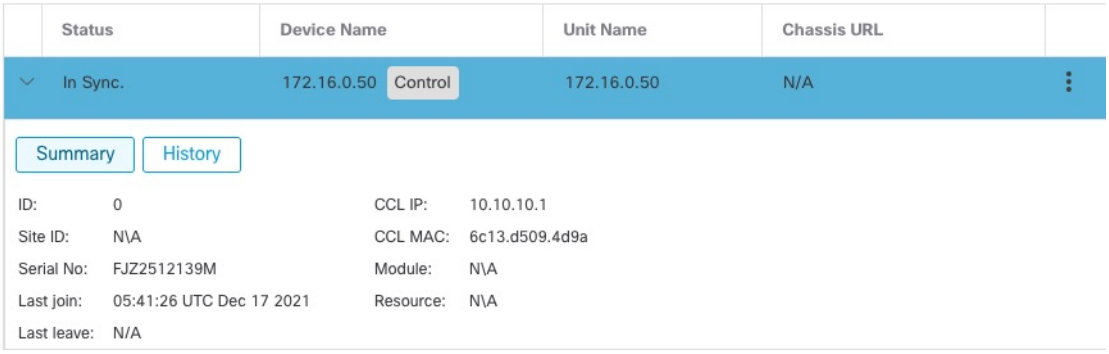
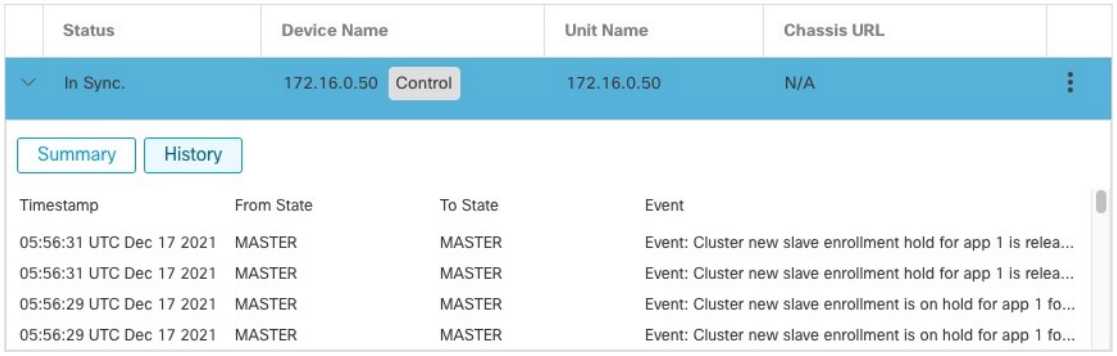


Figure 193: Node history



- **System** (⚙) > **Tasks** page.

  The **Tasks** page shows updates of the Cluster Registration task as each node registers.

- **Devices** > **Device Management** and then click **Add** > **Device** *cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member nodes, including the control node shown with its role next to the IP address. For nodes that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Cluster health monitor

When a Firewall Threat Defense is the control node of a cluster, the Cloud-Delivered Firewall Management Center collects various metrics periodically from the device metric data collector.

The cluster health monitor is comprised of these components:

- Overview dashboard: Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Cloud-Delivered Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.



Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard: Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.

- Member Performance dashboard: Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard: Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links: Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range: An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard: Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

View the cluster health monitor

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

Ensure you have created a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Follow these steps to view the cluster health monitor:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed cluster devices.

Step 3 Click on the cluster name to view the cluster health statistics.

The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution: Traffic and packet distribution across the cluster nodes.
- Member Performance: Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL: Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 View device monitor health and performance metrics in several predefined dashboards.

The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.

- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics that are related to the Snort process.
- ASP drops: Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. Refer to Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 60: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number
Packets	Packet distribution count in the cluster for every second.	number

Upgrade the cluster

This task upgrades a Firewall Threat Defense Virtual cluster to a new software version by updating the cloud instance template and performing a managed upgrade.

The cluster upgrade process involves updating both the cloud infrastructure template and the software package through the management center. The upgrade is performed one node at a time to maintain cluster availability.

Before you begin

- Before you upgrade a cluster in the public cloud, copy the target version image to your cloud image repository and update the image ID in the cluster deployment template (we actually recommend replacing the existing template with a modified copy). This ensures that after the upgrade, new instances — for example, instances launched during cluster scaling — will use the correct version. If the marketplace does not have the image you need, such as when the cluster has been patched, create a custom image from a snapshot of a standalone Firewall Threat Defense Virtual instance running the correct version, with no instance-specific (day 0) configurations.

Follow these steps to upgrade a Firewall Threat Defense Virtual cluster:

Procedure

Step 1 Upload the target image version to the cloud image storage.

Step 2 Update the cloud instance template of the cluster with the updated target image version.

a) Create a copy of the instance template with the target image version.

b) Attach the newly created template to cluster instance group.

Step 3 Upload the target image version upgrade package to the Cloud-Delivered Firewall Management Center.

Step 4 Perform a readiness check on the cluster that you want to upgrade.

Step 5 After a successful readiness check, initiate the installation of the upgrade package.

Step 6 The Cloud-Delivered Firewall Management Center upgrades the cluster nodes one at a time.

Step 7 The Cloud-Delivered Firewall Management Center displays a notification after successful upgrade of the cluster.

There is no change to the serial number and UUID of the instance after the upgrade.

Reference for clustering

This section includes more information about how clustering operates.

Threat Defense features and clustering

Some Firewall Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported features and clustering

Centralized features for clustering

These features are only supported on the control node, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member nodes to the control node over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control nodes before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control node.

For centralized features, if the control node fails, all connections are dropped, and you have to re-establish the connections on the new control node.

Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Cloud-Delivered Firewall Management Center GUI. For more information, refer to FlexConfig policies, on page 1309.

- These application inspections are supported:
  - DCERPC and ESMTP
  - NetBIOS and PPTP
  - RSH
  - SQLNET and SUNRPC
  - TFTP
  - XDMCP
- Static route monitoring
- Site-to-site VPN
- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Dynamic routing

Cisco Trustsec and clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

Connection settings and clustering

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

Dynamic routing and clustering

Dynamic routing and clustering are network configuration approaches that

- enable each node to run the routing protocol as a standalone router in Individual interface mode,
- allow routes to be learned by each node independently, and
- require separate router IDs for each node when communicating with external routers.

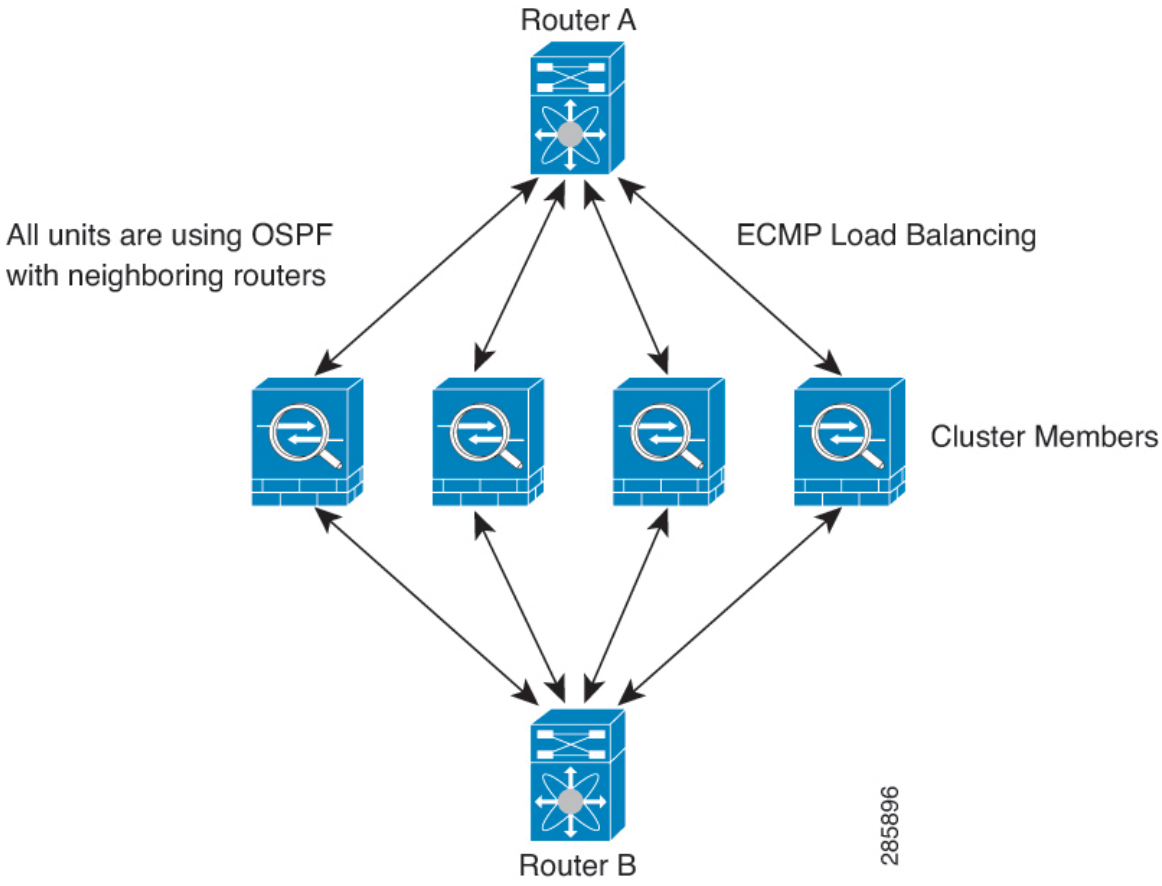
Routing behavior and load balancing

In Individual interface mode, ECMP is used to load balance traffic between multiple equal-cost paths. External routers can discover multiple paths to destinations through different cluster nodes.

Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when communicating with external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

Figure 194: Dynamic routing in individual interface mode



FTP and clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

NAT and clustering

Provides information about NAT behavior and limitations in clustering deployments.

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different Firewall Threat Defenses in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the Firewall Threat Defense that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- No Proxy ARP—For individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with an ASA that may no longer be in the cluster. The upstream router requires a static route or policy-based routing (PBR) with Object Tracking for the mapped addresses, which must point to the main cluster IP address.
- No interface PAT on an individual interface—Interface PAT is not supported for Individual interfaces.
- PAT with Port Block Allocation—See the following guidelines for this feature:
  - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
  - Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
  - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
  - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.
- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
  - FTP

- RSH
- SQLNET
- TFTP
- XDMCP
- SIP

- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

SIP inspection and clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and clustering

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

VPN and Clustering

Site-to-site VPN is a centralized feature; only the control node supports VPN connections.



Note Remote access VPN is not supported with clustering.

VPN functionality is limited to the control node and does not take advantage of the cluster high availability capabilities. If the control node fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control node is elected, you must reestablish the VPN connections.

For connections to an Individual interface when using PBR or ECMP, you must always connect to the Main cluster IP address, not a Local address.

VPN-related keys and certificates are replicated to all nodes.

Performance scaling factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

Throughput calculation

For example, if your model can handle approximately 10 Gbps of traffic when running alone, then for a cluster of 8 units, the maximum combined throughput will be approximately 80% of 80 Gbps (8 units x 10 Gbps): 64 Gbps.

Control node election

Summary

The key components involved in the control node election process are:

- Cluster nodes: Communicate over the cluster control link and participate in the election process
- Election requests: Broadcast every 3 seconds by nodes when clustering is enabled
- Priority values: Set between 1 and 100 where 1 is the highest priority
- Control node: The elected node that coordinates cluster operations

Workflow

These stages describe how cluster nodes elect a control node:

1. When clustering is enabled for a node (or when it first starts up with clustering already enabled), the node broadcasts an election request every 3 seconds.
2. Any other nodes with a higher priority respond to the election request. The priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a node does not receive a response from another node with a higher priority, then it becomes the control node.



Note If multiple nodes tie for the highest priority, the cluster node name and then the serial number is used to determine the control node.

4. If a node later joins the cluster with a higher priority, it does not automatically become the control node. The existing control node always remains as the control node unless it stops responding, at which point a new control node is elected.
5. In a "split brain" scenario when there are temporarily multiple control nodes, the node with highest priority retains the role while the other nodes return to data node roles.

Note You can manually force a node to become the control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node.

High availability within the cluster

Clustering provides high availability by monitoring node and interface health and by replicating connection states between nodes.

Node health monitoring

Each node periodically sends a broadcast heartbeat packet over the cluster control link. If the control node does not receive any heartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining nodes.

Split brain scenario

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role.

Interface monitoring

Each node monitors the link status of all named hardware interfaces in use, and reports status changes to the control node.

All physical interfaces are monitored; only named interfaces can be monitored. You can optionally disable monitoring per interface.

A node is removed from the cluster if its monitored interfaces fail. The node is removed after 500 milliseconds.

Status after failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The Firewall Threat Defense automatically tries to rejoin the cluster, depending on the failure event.

Note When the Firewall Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoining the cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firewall Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firewall Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firewall Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from Cloud-Delivered Firewall Management Center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.

Data path connection state replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 61: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—

Traffic	State Support	Notes
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the cluster manages connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.
- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not

used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

Port Address Translation Connections

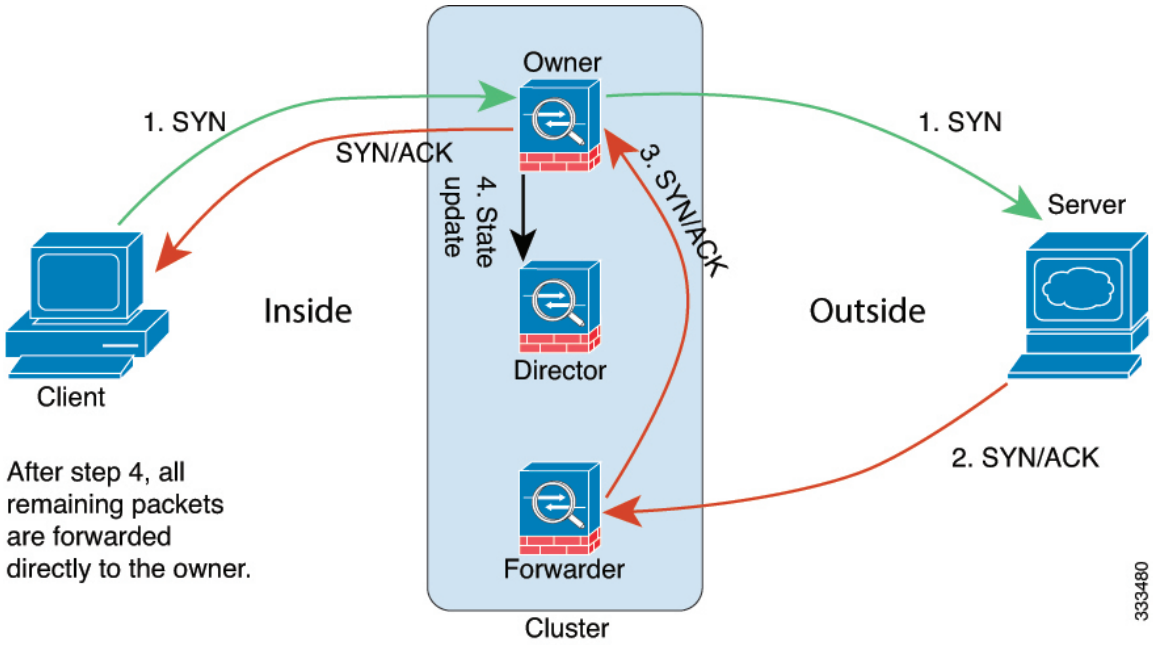
New connection ownership

When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. If a reverse flow arrives at a different node, it is redirected back to the original node.

Traffic redirection is not supported in this release. When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. All the subsequent packets for the same connection should arrive the same node. If any connection packets arrive at a different node, they will be dropped. If a reverse flow arrives at a different node, it will be dropped as well. For centralized features, if the connections do not arrive on the control node, they will be dropped.

Sample data flow for TCP

The following example shows the establishment of a new connection.

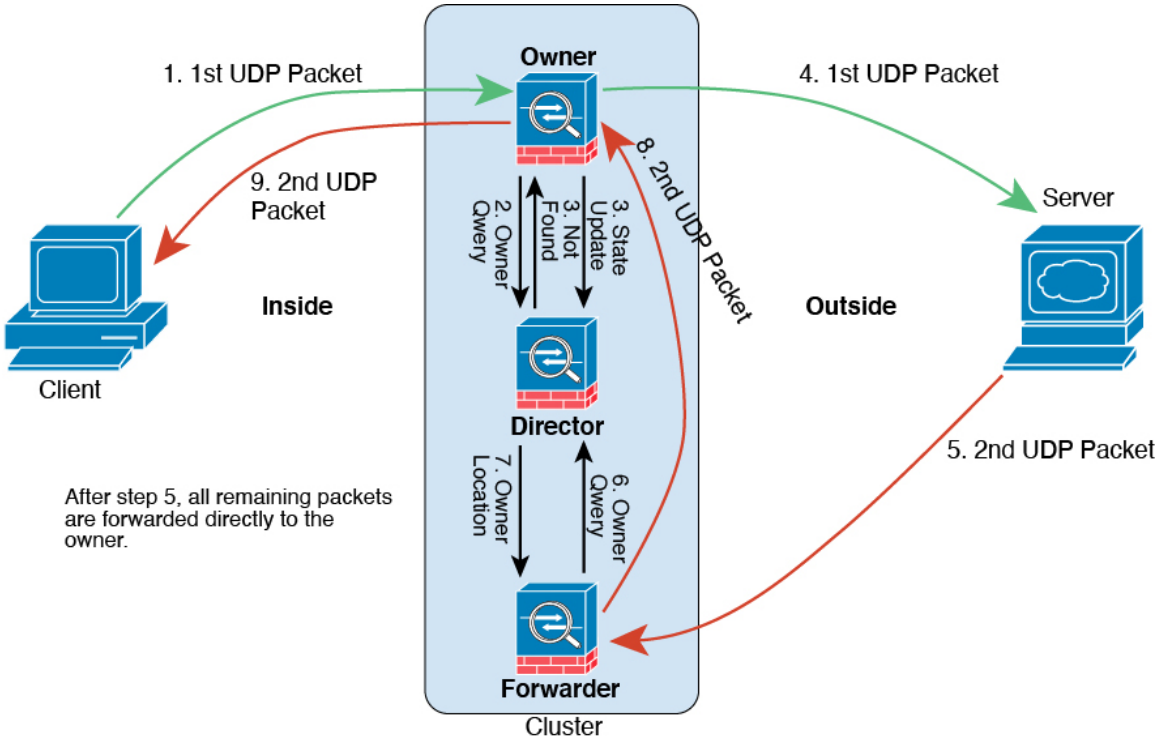


1. The SYN packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different Firewall Threat Defense (based on the load balancing method). This Firewall Threat Defense is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample data flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. *Figure 195: ICMP and UDP Data Flow*



The first UDP packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

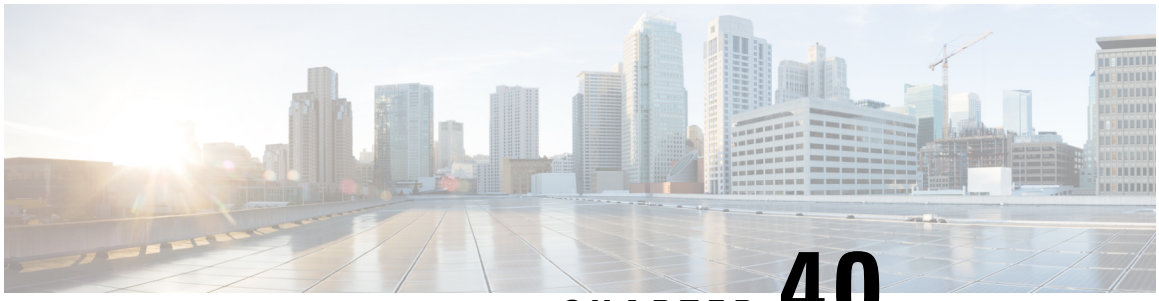
History for Threat Defense Virtual clustering in the public cloud

This reference provides the historical timeline and version information for Threat Defense Virtual clustering features introduced across different software releases and deployment platforms in public cloud environments.

Table 62: Feature history

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
MTU ping test from data node on node join	20241030	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size twice the cluster control link MTU. Previously, only the control node sent a ping. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again. Added/modified commands: **show cluster history**. See: Clustering for the Secure Firewall 3100/4200, Clustering for Threat Defense Virtual in a Private Cloud, Clustering for Threat Defense Virtual in a Public Cloud, Clustering for the Firepower 4100/9300
Cluster health monitor settings.	20221213	Any	You can now edit cluster health monitor settings. New/Modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard.	20221213	Any	You can now view cluster health on the cluster health monitor dashboard. New/Modified screens: **System** > **Health** > **Monitor**
Clustering for the Firewall Threat Defense Virtual in Azure.	20221213	7.3.0	You can now configure clustering for up to 16 nodes the Firewall Threat Defense Virtual in Azure for the Azure Gateway Load Balancer or for external load balancers. New/modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense Virtual in Azure

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Clustering for the Firewall Threat Defense Virtual in the Public Cloud (Amazon Web Services and Google Cloud Platform).	Any	7.2.0	The Firewall Threat Defense Virtual supports Individual interface clustering for up to 16 nodes in the public cloud (AWS and GCP). New/Modified screens: • **Devices** > **Device Management** > **Add Device** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense Virtual in AWS and GCP



CHAPTER 40

Clustering: Firepower 4100/9300

Clustering lets you group multiple Firewall Threat Defense nodes together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices.

Note Some features are not supported when using clustering. See Unsupported features with clustering, on page 977.

- About clustering on the Firepower 4100/9300 chassis, on page 935
- Licenses for clustering, on page 939
- Requirements and prerequisites for clustering, on page 940
- Clustering guidelines and limitations, on page 944
- Configure clustering, on page 947
- FXOS: Remove a cluster node, on page 962
- Cloud-Delivered Firewall Management Center: Manage cluster members, on page 964
- Cloud-Delivered Firewall Management Center: Monitor the cluster, on page 969
- Examples for Clustering, on page 976
- Reference for clustering, on page 977
- History for clustering, on page 989

About clustering on the Firepower 4100/9300 chassis

When you deploy a cluster on the Firepower 4100/9300 chassis, it does the following:

- For native instance clustering: Creates a *cluster-control link* (by default, port-channel 48) for node-to-node communication.

  For multi-instance clustering: You should pre-configure subinterfaces on one or more cluster-type EtherChannels; each instance needs its own cluster control link.

  For a cluster isolated to security modules within one Firepower 9300 chassis, this link utilizes the Firepower 9300 backplane for cluster communications.

  For clustering with multiple chassis, you need to manually assign physical interface(s) to this EtherChannel for communications between chassis.

- Creates the cluster bootstrap configuration within the application.

  When you deploy the cluster, the chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings.

- Assigns data interfaces to the cluster as *Spanned* interfaces.

  For a cluster isolated to security modules within one Firepower 9300 chassis, spanned interfaces are not limited to EtherChannels, like it is for clustering with multiple chassis. The Firepower 9300 supervisor uses EtherChannel technology internally to load-balance traffic to multiple modules on a shared interface, so any data interface type works for Spanned mode. For clustering with multiple chassis, you must use Spanned EtherChannels for all data interfaces.



Note Individual interfaces are not supported, with the exception of a management interface.

- Assigns a management interface to all units in the cluster.

Bootstrap configuration

A bootstrap configuration is a minimal cluster configuration that includes cluster name, cluster control link interface, and other cluster settings.

When you deploy the cluster, the Firepower 4100/9300 chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings.

Cluster members

Cluster members are security devices that work together to accomplish the sharing of the security policy and traffic flows.

Cluster member roles and management

One member of the cluster is the **control** unit. The control unit is determined automatically. All other members are **data** units.

You must perform all configuration on the control unit only; the configuration is then replicated to the data units.

Some features do not scale in a cluster, and the control unit handles all traffic for those features. .

Cluster control link

Cluster control link traffic includes both control and data traffic.

For native instance clustering: The cluster control link is automatically created using the Port-channel 48 interface.

For multi-instance clustering: You should pre-configure subinterfaces on one or more cluster-type EtherChannels; each instance needs its own cluster control link.

For a cluster isolated to security modules within one Firepower 9300 chassis, this interface has no member interfaces. This Cluster type EtherChannel utilizes the Firepower 9300 backplane for cluster communications. For clustering with multiple chassis, you must add one or more interfaces to the EtherChannel.

For a cluster with two chassis, do not directly connect the cluster control link from one chassis to the other chassis. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit.

Size the cluster control link

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features.

For example:

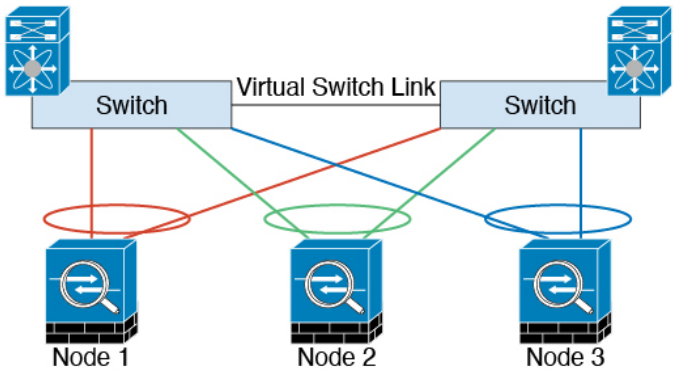
- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.
- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.

Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster control link redundancy

The diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS), Virtual Port Channel (vPC), StackWise, or StackWise Virtual environment. All links in the EtherChannel are active. When the switch is part of a redundant system, then you can connect firewall interfaces within the same EtherChannel to separate switches in the redundant system. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a Spanned EtherChannel.



Cluster control link reliability for inter-chassis clustering

To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

Cluster control link network

The Firepower 4100/9300 chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: 127.2.*chassis_id*.*slot_id*. For multi-instance clusters, which typically use different VLAN subinterfaces of the same EtherChannel, the same IP address can be used for different clusters because of VLAN separation. The cluster control link network cannot include any routers between units; only Layer 2 switching is allowed.

Management network

We recommend connecting all units to a single management network. This network is separate from the cluster control link.

Management interface

You must assign a Management type interface to the cluster. This interface is a special individual interface as opposed to a Spanned interface. The management interface lets you connect directly to each unit. This Management logical interface is separate from the other interfaces on the device. It is used to set up and register the device to the Cloud-Delivered Firewall Management Center. It uses its own local authentication, IP address, and static routing.

Management interface configuration

Each cluster member uses a separate IP address on the management network that you set as part of the bootstrap configuration.

Cluster interfaces

For a cluster isolated to security modules within one Firepower 9300 chassis, you can assign either physical interfaces or EtherChannels (also known as port channels) to the cluster. The assigned interfaces are Spanned interfaces that load-balance traffic across all cluster members.

For clustering with multiple chassis, you can only assign data EtherChannels to the cluster. These Spanned EtherChannels include the same member interfaces on each chassis. On the upstream switch, all these interfaces are part of a single EtherChannel, so the switch does not recognize that it is connected to multiple devices.

You can use regular firewall interfaces or IPS-only interfaces (inline sets or passive interfaces).

Individual interfaces are not supported except for a management interface.

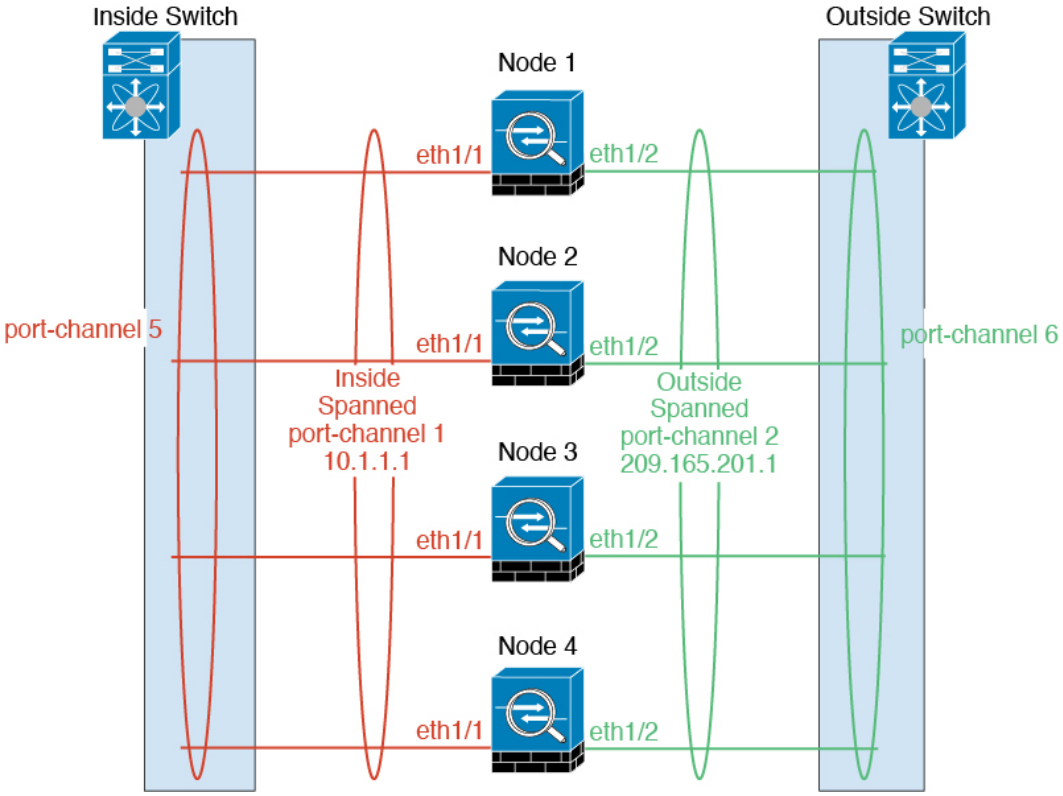
Spanned EtherChannels

You can group one or more interfaces per chassis into an EtherChannel that spans all chassis in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel.

For regular firewall interfaces:A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address is assigned to the BVI, not to the bridge group member interface.

The EtherChannel inherently provides load balancing as part of basic operation.

For multi-instance clusters, each cluster requires dedicated data EtherChannels; you cannot use shared interfaces or VLAN subinterfaces.



Configuration replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Licenses for clustering

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

Note APEX, PLUS, and VPN-Only licenses are not supported for clustering. Before you create a cluster, remove unsupported license assignments from the devices that will become cluster nodes.

When you add a cluster node to the Cloud-Delivered Firewall Management Center, you can specify the feature licenses you want to use for the cluster. You can modify licenses for the cluster in the **Devices** > **Device Management** and the choose the **Add**, **Cluster** and then **License** area.

Note If you add the cluster before the Cloud-Delivered Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Cloud-Delivered Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and prerequisites for clustering

This topic provides the necessary prerequisites and hardware support details for configuring clustering on supported security appliances.

Cluster model support

The Firewall Threat Defense supports clustering on the following models:

- Firepower 9300— You can include up to 16 nodes in the cluster. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules. Supports clustering with multiple chassis and clustering isolated to security modules within one chassis.
- Firepower 4100—Supported for up to 16 nodes using clustering with multiple chassis.

User roles

The following user roles are required to manage clustering configurations:

- Admin
- Access Admin
- Network Admin

Clustering hardware and software requirements

All chassis in a cluster:

- Native instance clustering—For the Firepower 4100: All chassis must be the same model. For the Firepower 9300: All security modules must be the same type. For example, if you use clustering, all modules in the Firepower 9300 must be SM-40s. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots.

- Container instance clustering—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.



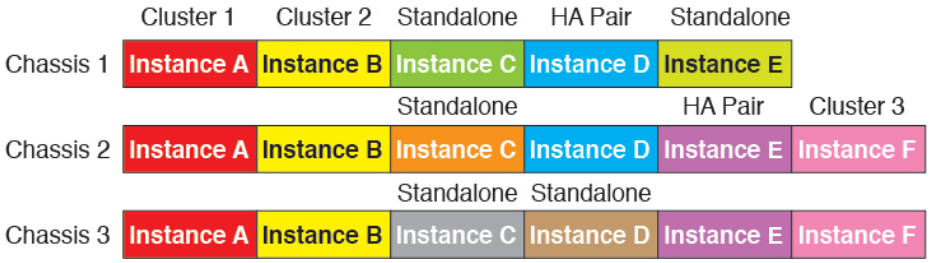
- Must run the identical FXOS and application software except at the time of an image upgrade. Mismatched software versions can lead to poor performance, so be sure to upgrade all nodes in the same maintenance window.
- Must include the same interface configuration for interfaces you assign to the cluster, such as the same Management interface, EtherChannels, active interfaces, speed and duplex, and so on. You can use different network module types on the chassis as long as the capacity matches for the same interface IDs and interfaces can successfully bundle in the same spanned EtherChannel. Note that all data interfaces must be EtherChannels in clusters with multiple chassis. If you change the interfaces in FXOS after you enable clustering (by adding or removing interface modules, or configuring EtherChannels, for example), then perform the same changes on each chassis, starting with the data nodes, and ending with the control node.
- Must use the same NTP server. For Firewall Threat Defense, the Cloud-Delivered Firewall Management Center must also use the same NTP server. Do not set the time manually.

Multi-instance clustering requirements

- No intra-security-module/engine clustering—For a given cluster, you can only use a single container instance per security module/engine. You cannot add 2 container instances to the same cluster if they are running on the same module.



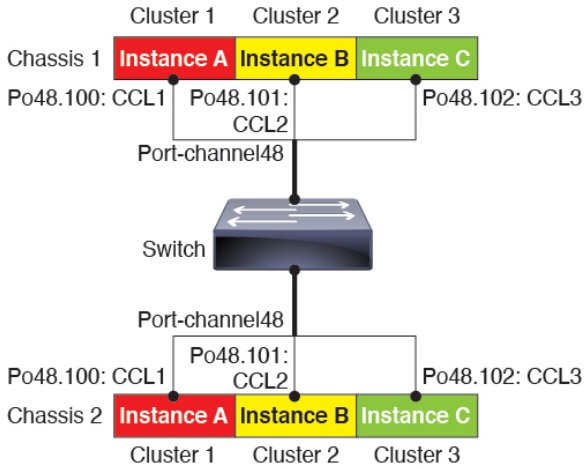
- Mix and match clusters and standalone instances—Not all container instances on a security module/engine need to belong to a cluster. You can use some instances as standalone or High Availability nodes. You can also create multiple clusters using separate instances on the same security module/engine.



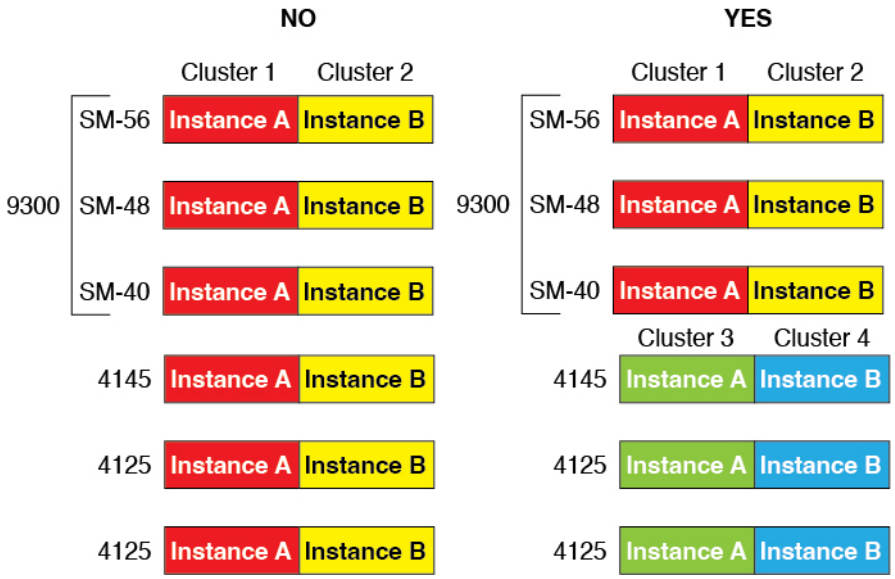
- All 3 modules in a Firepower 9300 must belong to the cluster—For the Firepower 9300, a cluster requires a single container instance on all 3 modules. You cannot create a cluster using instances on module 1 and 2, and then use a native instance on module 3, or example.



- Match resource profiles—We recommend that each node in the cluster use the same resource profile attributes; however, mismatched resources are allowed when changing cluster nodes to a different resource profile, or when using different models.
- Dedicated cluster control link—For clusters with multiple chassis, each cluster needs a dedicated cluster control link. For example, each cluster can use a separate subinterface on the same cluster-type EtherChannel, or use separate EtherChannels.



- No shared interfaces—Shared-type interfaces are not supported with clustering. However, the same Management and Eventing interfaces can used by multiple clusters.
- No subinterfaces—A multi-instance cluster cannot use FXOS-defined VLAN subinterfaces. An exception is made for the cluster control link, which can use a subinterface of the Cluster EtherChannel.
- Mix chassis models—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.



- Maximum 6 nodes—You can use up to six container instances in a cluster.

Switch requirements

- Be sure to complete the switch configuration and successfully connect all the EtherChannels from the chassis to the switch(es) before you configure clustering on the Firepower 4100/9300 chassis.

- For supported switch characteristics, see Cisco FXOS Compatibility.

Clustering guidelines and limitations

Clustering is a deployment method that groups multiple security devices into a single logical unit to provide high availability and increased throughput.

Switches for Clustering

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. In addition, we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.
- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.
- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **src-dst-mixed-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.
- Switches on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.
- Port-channel bundling downtime should not exceed the configured keepalive interval.
- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the cluster device to fixed:

  router(config)# **port-channel** *id* **hash-distribution fixed**

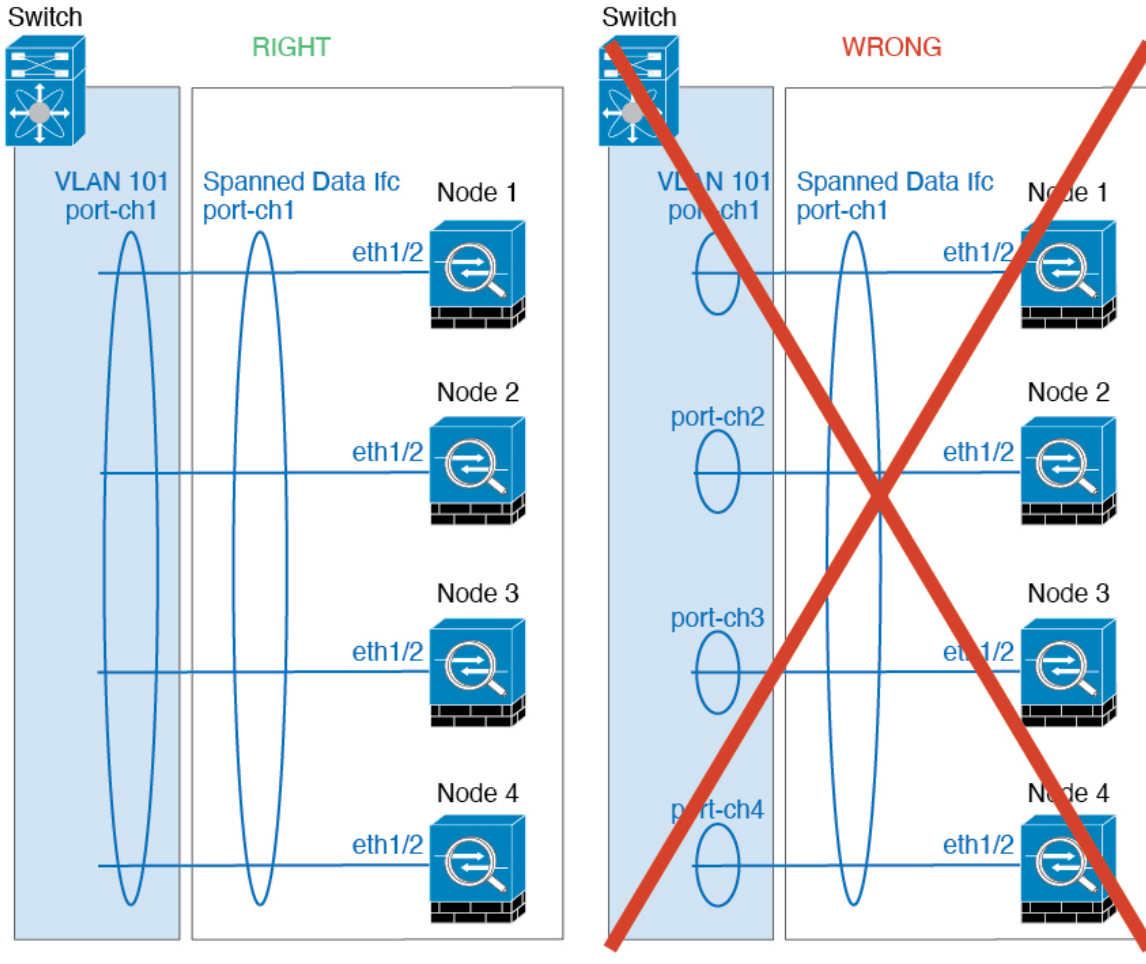
  Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

- Firepower 4100/9300 clusters support LACP graceful convergence. So you can leave LACP graceful convergence enabled on connected Cisco Nexus switches.
- When you see slow bundling of a Spanned EtherChannel on the switch, you can enable LACP rate fast for an individual interface on the switch. FXOS EtherChannels have the LACP rate set to fast by default.

Note that some switches, such as the Nexus series, do not support LACP rate fast when performing in-service software upgrades (ISSUs), so we do not recommend using ISSUs with clustering.

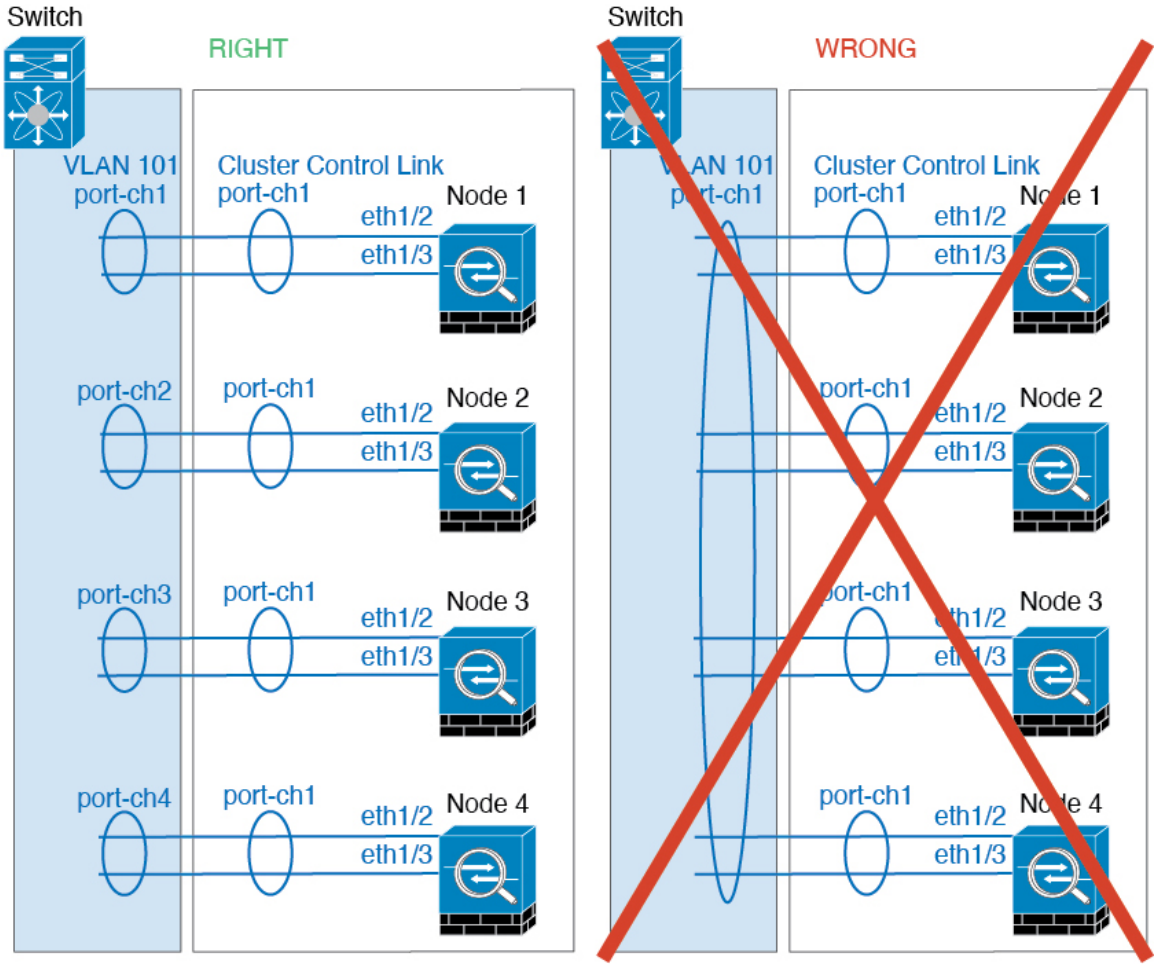
EtherChannels for Clustering

- In Catalyst 3750-X Cisco IOS software versions earlier than 15.1(1)S2, the cluster unit did not support connecting an EtherChannel to a switch stack. With default switch settings, if the cluster unit EtherChannel is connected cross stack, and if the control unit switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.
- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.
  - Spanned EtherChannels—For cluster unit *Spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.



  - Device-local EtherChannels—For cluster unit *Device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels

on the switch; do not combine multiple cluster unit EtherChannels into one EtherChannel on the switch.



Additional Guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Firepower 4100/9300 chassis or the switch, adding an additional switch to form a VSS, vPC, StackWise, or StackWise Virtual) you should disable the health check feature, and also disable interface monitoring for the disabled interfaces . When the topology change is complete, and the configuration change is synced to all units, you can re-enable the health check feature.
- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang connections; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel interface, when the syslog server port is down, and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the cluster. These messages can result in some units of the cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.

- We recommend connecting EtherChannels to a VSS, vPC, StackWise, or StackWise Virtual for redundancy.
- Within a chassis, you cannot cluster some security modules and run other security modules in standalone mode; you must include all security modules in the cluster.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Defaults

- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is set to unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is set to 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure clustering

You can easily deploy the cluster from the Firepower 4100/9300 supervisor. All initial configuration is automatically generated for each unit. You can then add the units to the Cloud-Delivered Firewall Management Center and group them into a cluster.

FXOS: Add a Firewall Threat Defense cluster

A cluster is a deployment configuration that allows you to group security modules to increase capacity and redundancy.

Cluster Deployment Modes

In native mode: You can add a cluster to a single Firepower 9300 chassis that is isolated to security modules within the chassis, or you can use multiple chassis.

In multi-instance mode: You can add one or more clusters to a single Firepower 9300 chassis that are isolated to security modules within the chassis (you must include an instance on each module), or add one or more clusters on multiple chassis.

For clusters on multiple chassis, you must configure each chassis separately. Add the cluster on one chassis; you can then

Create a Firewall Threat Defense cluster

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit.

For clustering on multiple chassis, you must configure each chassis separately. Deploy the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, or for container instances, a container instance in each slot, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.
- For container instances, if you do not want to use the default profile, add a resource profile according to Add a resource profile for container instances.
- For container instances, before you can install a container instance for the first time, you must reinitialize the security module/engine so that the disk has the correct formatting. An existing logical device will be deleted and then reinstalled as a new device, losing any local application configuration. If you are replacing a native instance with container instances, you will need to delete the native instance in any case. You cannot automatically migrate a native instance to a container instance.
- Gather the following information:
  - Management interface ID, IP addresses, and network mask
  - Gateway IP address
  - Firewall Management Center IP address and/or NAT ID of your choosing
  - DNS server IP address
  - Firewall Threat Defense hostname and domain name

Procedure

Step 1 Configure interfaces.

a) Add at least one Data type interface or EtherChannel (also known as a port-channel) before you deploy the cluster. See Add an EtherChannel (port channel) or Configure a physical interface.

For clustering on multiple chassis, all data interfaces must be Spanned EtherChannels with at least one member interface. Add the same EtherChannels on each chassis. Combine the member interfaces from all cluster units into a single EtherChannel on the switch. See Clustering guidelines and limitations, on page 944 for more information about EtherChannels.

For multi-instance clustering, you cannot use FXOS-defined VLAN subinterfaces or data-sharing interfaces in the cluster. Only application-defined subinterfaces are supported. See FXOS interfaces and application interfaces for more information.

b) Add a Management type interface or EtherChannel. See Add an EtherChannel (port channel) or Configure a physical interface.

The management interface is required. Note that this management interface is not the same as the chassis management interface that is used only for chassis management (in FXOS, you might see the chassis management interface displayed as MGMT, management0, or other similar names).

For clustering on multiple chassis, add the same Management interface on each chassis.

For multi-instance clustering, you can share the same management interface across multiple clusters on the same chassis, or with standalone instances.

c) For clustering on multiple chassis, add a member interface to the cluster control link EtherChannel (by default, port-channel 48). See Add an EtherChannel (port channel).

Do not add a member interface for a cluster isolated to security modules within one Firepower 9300 chassis. If you add a member, the chassis assumes this cluster will be using multiple chassis, and will only allow you to use Spanned EtherChannels, for example.

Add the same member interfaces on each chassis. The cluster control link is a device-local EtherChannel on each chassis. Use separate EtherChannels on the switch per device. See Clustering guidelines and limitations, on page 944 for more information about EtherChannels.

For multi-instance clustering, you can create additional Cluster type EtherChannels. Unlike the Management interface, the cluster control link is not sharable across multiple devices, so you will need a Cluster interface for each cluster. However, we recommend using VLAN subinterfaces instead of multiple EtherChannels; see the next step to add a VLAN subinterface to the Cluster interface.

d) For multi-instance clustering, add VLAN subinterfaces to the cluster EtherChannel so you have a subinterface for each cluster. See Add a VLAN subinterface for container instances.

If you add subinterfaces to a Cluster interface, you cannot use that interface for a native cluster.

e) (Optional) Add an eventing interface. See Add an EtherChannel (port channel) or Configure a physical interface.

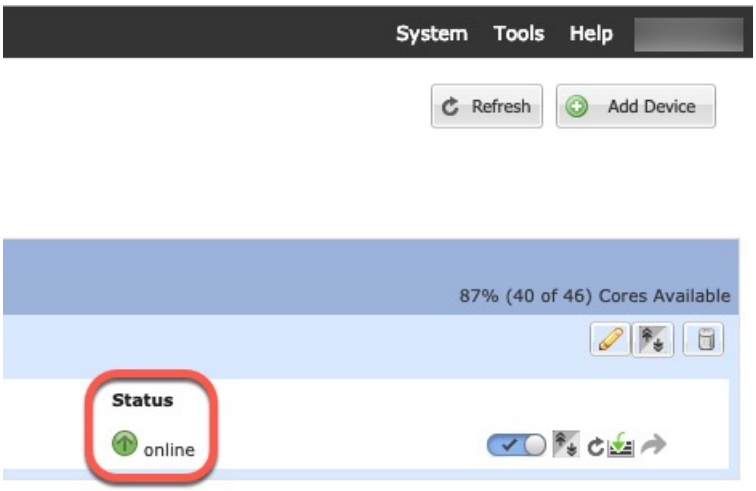
This interface is a secondary management interface for the Firewall Threat Defense devices. To use this interface, you must configure its IP address and other parameters at the Firewall Threat Defense CLI. For example, you can separate management traffic from events (such as web events). See the **configure network** commands in the Firewall Threat Defense command reference.

For clustering on multiple chassis, add the same eventing interface on each chassis.

Step 2 Choose **Logical Devices**.

Step 3 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can add the remaining cluster chassis, or for a cluster isolated to security modules within one Firepower 9300 chassis, start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



Add more cluster nodes

Add or replace the Firewall Threat Defense cluster node in an existing cluster. When you add a new cluster node in FXOS, the Cloud-Delivered Firewall Management Center adds the node automatically.

Note The FXOS steps in this procedure only apply to adding a new chassis ; if you are adding a new module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

- In the case of a replacement, you must delete the old cluster node from the Cloud-Delivered Firewall Management Center . When you replace it with a new node, it is considered to be a new device on the Cloud-Delivered Firewall Management Center .
- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.

What to do next

1. To add another chassis to the cluster, repeat the procedure in Create a Firewall Threat Defense cluster , on page 947 except you must configure the following settings to be unique; otherwise, use the same configuration for both chassis.
   - Chassis ID
   - Management IP addresses

   Also be sure to set the startup version to the currently running version on the cluster nodes.
2. In the Cloud-Delivered Firewall Management Center , choose **Devices** > **Device Management** , and select the **Add** > **Device** to add the new logical device.
3. Choose **Add** > **Cluster** .

4. Choose the current control unit from the drop-down list.

   When you choose a control unit that is already in a cluster, then the existing cluster name is auto-filled, and all eligible data units are added to the box, including the new unit you just added to the Cloud-Delivered Firewall Management Center .

5. Click **Add** , and then **Deploy** . The cluster is updated to include the new member.

Cloud-Delivered Firewall Management Center: Add a cluster

Add one of the cluster units as a new device to the Cloud-Delivered Firewall Management Center; the Cloud-Delivered Firewall Management Center auto-detects all other cluster members.

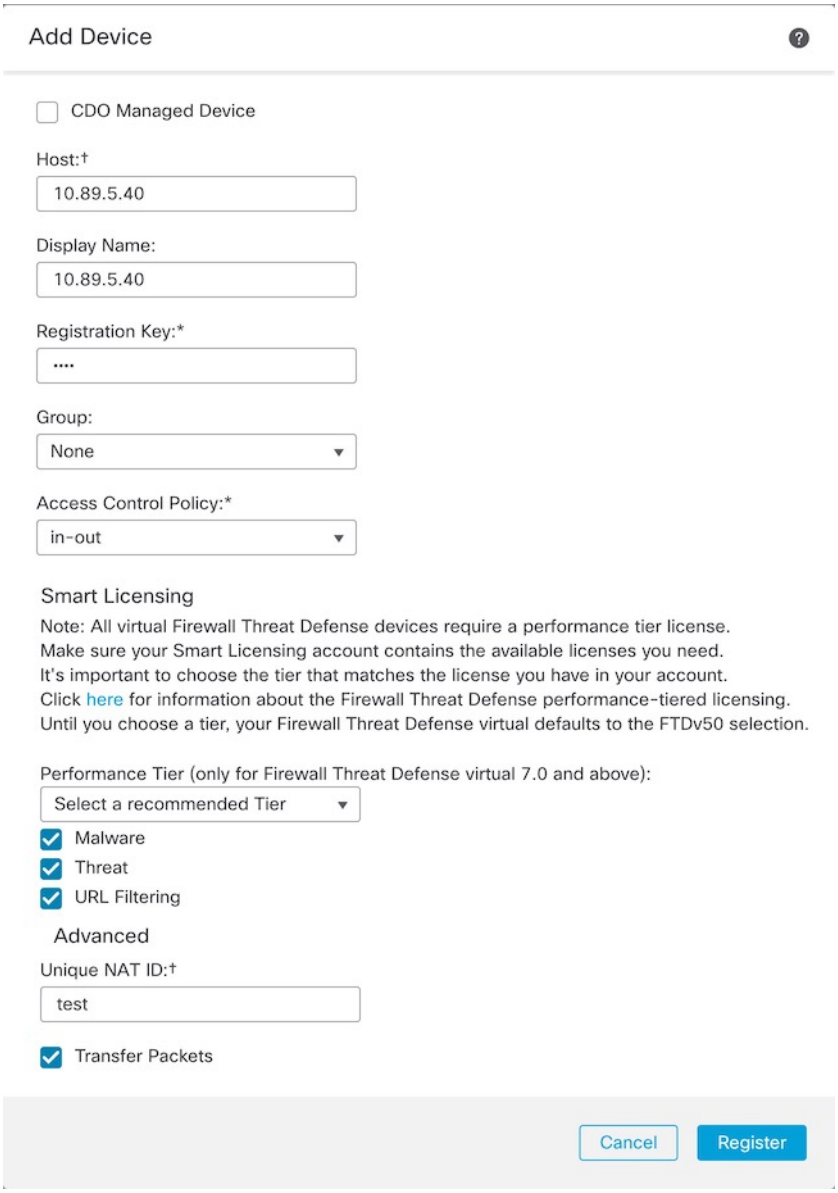
Before you begin

- All cluster units must be in a successfully-formed cluster on FXOS prior to adding the cluster to the Cloud-Delivered Firewall Management Center. You should also check which unit is the control unit. Refer to the Firewall Chassis Manager **Logical Devices** screen or use the Firewall Threat Defense **show cluster info** command.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and then choose **Add** > **Add Device** to add the control unit using the unit's management IP address you assigned when you deployed the cluster.

Figure 196: Add Device



a) In the **Host** field, enter the IP address or hostname of the control unit.

   We recommend adding the control unit for the best performance, but you can add any unit of the cluster.

   If you used a NAT ID during device setup, you may not need to enter this field.

b) In the **Display Name** field, enter a name for the control unit as you want it to display in the Cloud-Delivered Firewall Management Center.

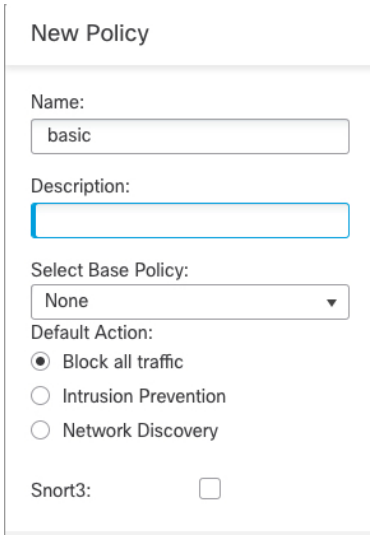
   This display name is not for the cluster; it is only for the control unit you are adding. You can later change the name of other cluster members and the cluster display name.

c) In the **Registration Key** field, enter the same registration key that you used when you deployed the cluster in FXOS. The registration key is a one-time-use shared secret.

d) (Optional) Add the device to a device **Group**.

e) Choose an initial **Access Control Policy** to deploy to the device upon registration, or create a new policy.

If you create a new policy, you create a basic policy only. You can later customize the policy as needed.



f) Choose licenses to apply to the device.

g) If you used a NAT ID during device setup, expand the **Advanced** section and enter the same NAT ID in the **Unique NAT ID** field.

h) Check the **Transfer Packets** check box to allow the device to transfer packets to the Cloud-Delivered Firewall Management Center.

This option is enabled by default. When events like IPS or Snort are triggered with this option enabled, the device sends event metadata information and packet data to the Cloud-Delivered Firewall Management Center for inspection. If you disable it, only event information will be sent to the Cloud-Delivered Firewall Management Center but packet data is not sent.

i) Click **Register**.

The Cloud-Delivered Firewall Management Center identifies and registers the control unit, and then registers all data units. If the control unit does not successfully register, then the cluster is not added. A registration failure can occur if the cluster was not up on the chassis, or because of other connectivity issues. In this case, we recommend that you try re-adding the cluster unit.

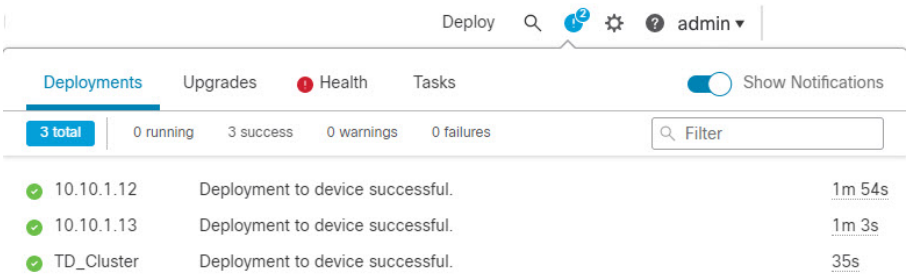
The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster units.



A unit that is currently registering shows the loading icon.



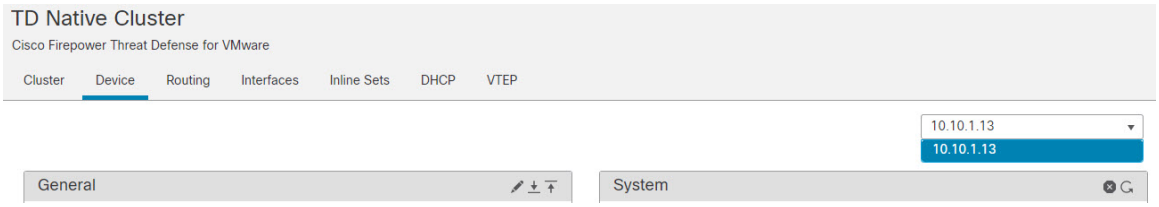
You can monitor cluster unit registration by clicking the **Notifications** icon and choosing **Tasks**. The Cloud-Delivered Firewall Management Center updates the Cluster Registration task as each unit registers. If any units fail to register, see Reconcile cluster members, on page 969.



Step 2 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

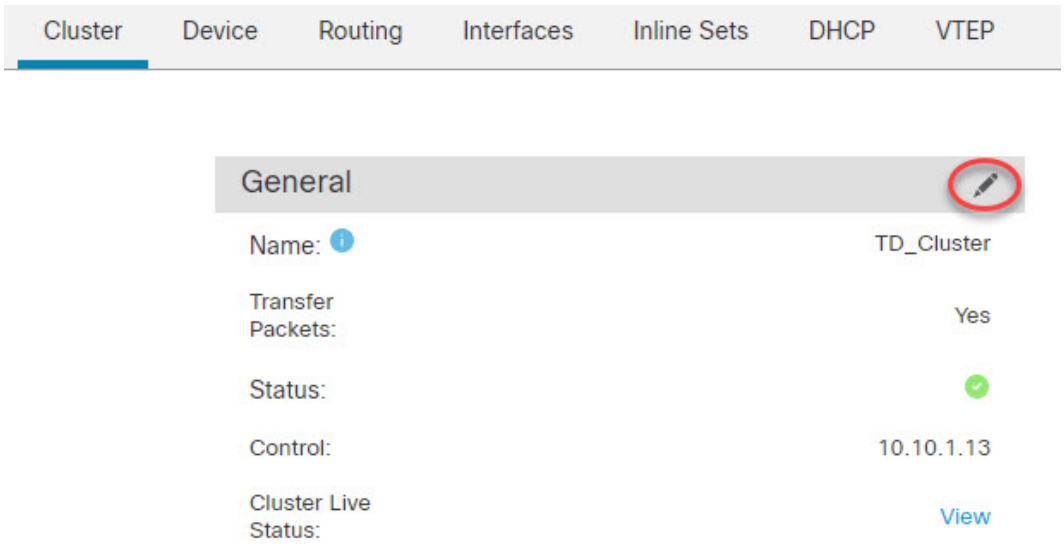
Most configuration can be applied to the cluster as a whole, and not member units in the cluster. For example, you can change the display name per unit, but you can only configure interfaces for the whole cluster.

Step 3 On the **Devices** > **Device Management**and choose the **Add**, **Cluster** screen, you see **General**, **License**, **System**, and **Health** settings.

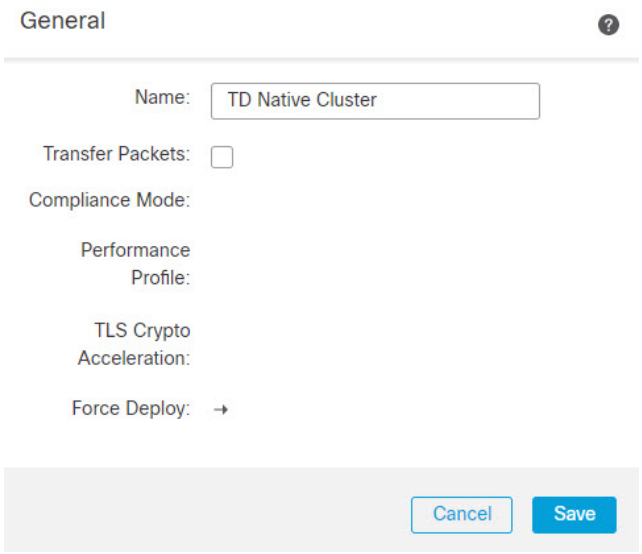


See the following cluster-specific items:

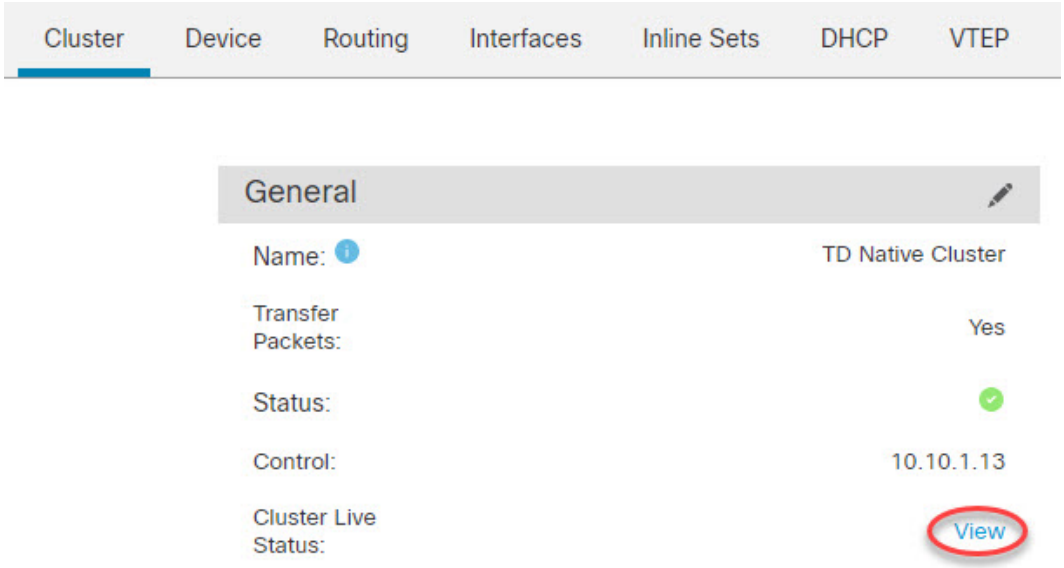
- **General > Name**—Change the cluster display name by clicking the **Edit** (✎).



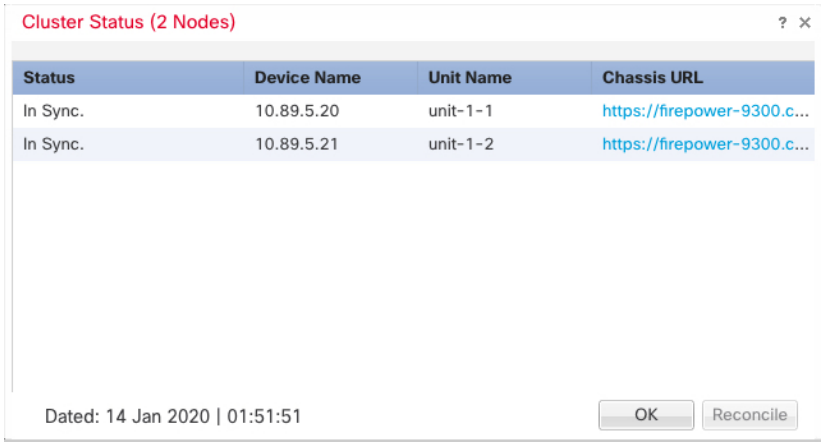
Then set the **Name** field.



- **General > View cluster status**—Click the **View cluster status** link to open the **Cluster Status** dialog box.



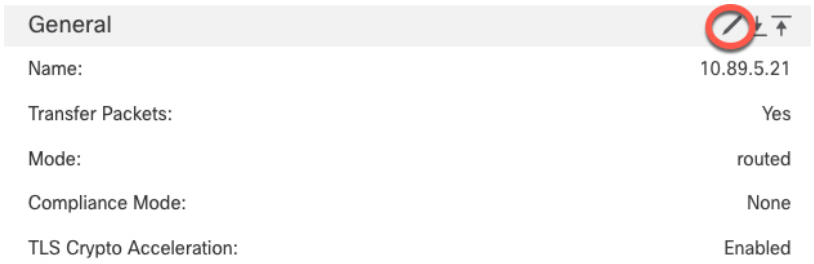
The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile**.



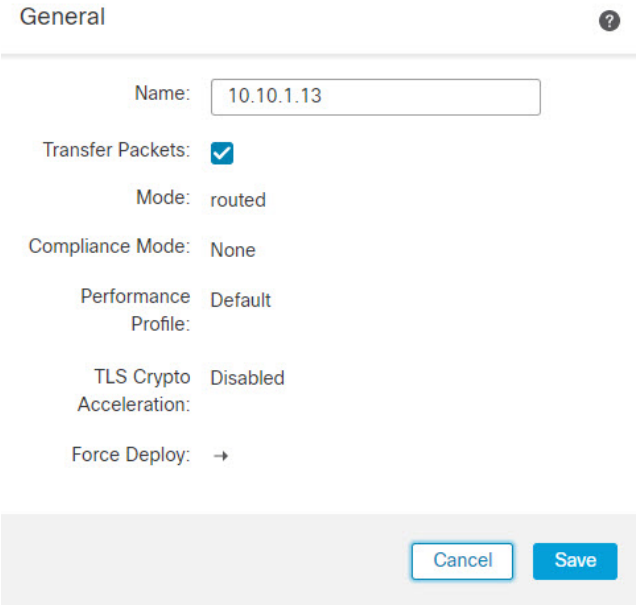
- **License**—Click **Edit** (✎) to set license entitlements.

Step 4 On the **Devices** > **Device Management** and then choose **Add Device**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.

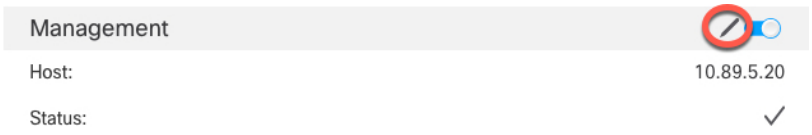
- **General > Name**—Change the cluster member display name by clicking the **Edit** (✎).



Then set the **Name** field.



- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the Cloud-Delivered Firewall Management Center so that it can reach the device on the network; edit the **Host** address in the **Management** area.



Cloud-Delivered Firewall Management Center: Configure cluster, data interfaces

This procedure configures basic parameters for each data interface that you assigned to the cluster when you deployed it in FXOS. For clustering on multiple chassis, data interfaces are always Spanned EtherChannel interfaces. For the cluster control link interface for a cluster isolated to security modules within one Firepower 9300 chassis, you must increase the MTU from the default.

Note When using Spanned EtherChannels for clustering on multiple chassis, the port-channel interface will not come up until clustering is fully enabled. This requirement prevents traffic from being forwarded to a unit that is not an active unit in the cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) next to the cluster.

Step 2 Click **Interfaces**.

Step 3 Configure the cluster control link.

For clustering on multiple chassis, set the cluster control link MTU to be at least 100 bytes higher than the highest MTU of the data interfaces. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. We suggest setting the MTU to the maximum of 9184; the minimum value is 1400 bytes. For example, because the maximum MTU is 9184, then the highest data interface MTU can be 9084, while the cluster control link can be set to 9184.

When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the initial ping fails, the node tries a ping using a smaller packet size (the MTU divided by 2, then by 4, then by 8) until a ping succeeds. A notification is generated so you can fix the MTU mismatch on connecting switches and try again.

For native clusters: The cluster control link interface is Port-Channel48 by default. If you don't know which interface is the cluster control link, check the FXOS configuration for chassis for the Cluster-type interface assigned to the cluster.

a) Click **Edit** (✎) for the cluster control link interface.
b) On the **General** page, in the **MTU** field, enter a value between 1400 and 9184but not between 2561 and 8362. Due to block pool handling, this MTU size is not optimal for system operation. We suggest using the maximum, 9184.
c) Click **OK**.

Step 4 Configure data interfaces.

a) (Optional) For regular firewall interfaces, configure VLAN subinterfaces on the data interface. The rest of this procedure applies to the subinterfaces. See Add a subinterface, on page 1060.
b) Click **Edit** (✎) for the data interface.
c) Configure the name and other parameters. For regular firewall interfaces, see Configure routed mode interfaces, on page 1080 or, for transparent mode, Bridge group interfaces, on page 1086. For IPS-only interfaces, see Inline Sets and Passive Interfaces, on page 1119..

Note
If the cluster control link interface MTU is not at least 100 bytes higher than the data interface MTU, you will see an error that you must reduce the MTU of the data interface. See, Step 3, on page 958 to increase the cluster control link MTU, after which you can continue configuring the data interfaces.

d) For clustering on multiple chassis, set a unique. manual global MAC address for the EtherChannel. Click **Advanced**, and in the **Active Mac Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

Do not set the **Standby Mac Address**; it is ignored.

You must configure a unique MAC address not currently in use on your network for a Spanned EtherChannel to avoid potential network connectivity problems. With a manually-configured MAC address, the MAC address stays with the current control unit. If you do not configure a MAC address, then if the control unit changes, the new control unit uses a new MAC address for the interface, which can cause a temporary network outage.

e) Click **OK**. Repeat the above steps for other data interfaces.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Cloud-Delivered Firewall Management Center: Configure Cluster Health Monitor Settings

This task allows you to modify cluster health monitor settings to control how cluster nodes monitor system health and automatically rejoin the cluster after failures.

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings for monitoring cluster node health. You can monitor any port-channel ID, any single physical interface ID, and the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link. It is always monitored.

When a topology change occurs (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature. Also, disable interface monitoring for any interfaces that are disabled. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Procedure

Step 1 Choose **Devices** > **Device Management**.

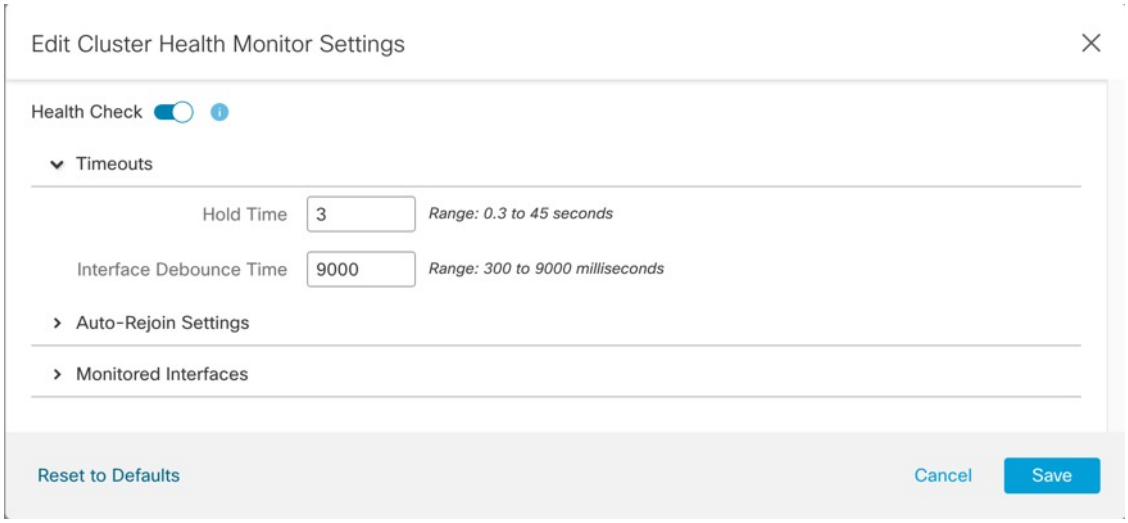
Step 2 Next to the cluster you want to modify, click **Edit** ( ).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** ( ).

Step 5 Disable the system health check by clicking the **Health Check** slider.

Figure 197: Disable the system health check



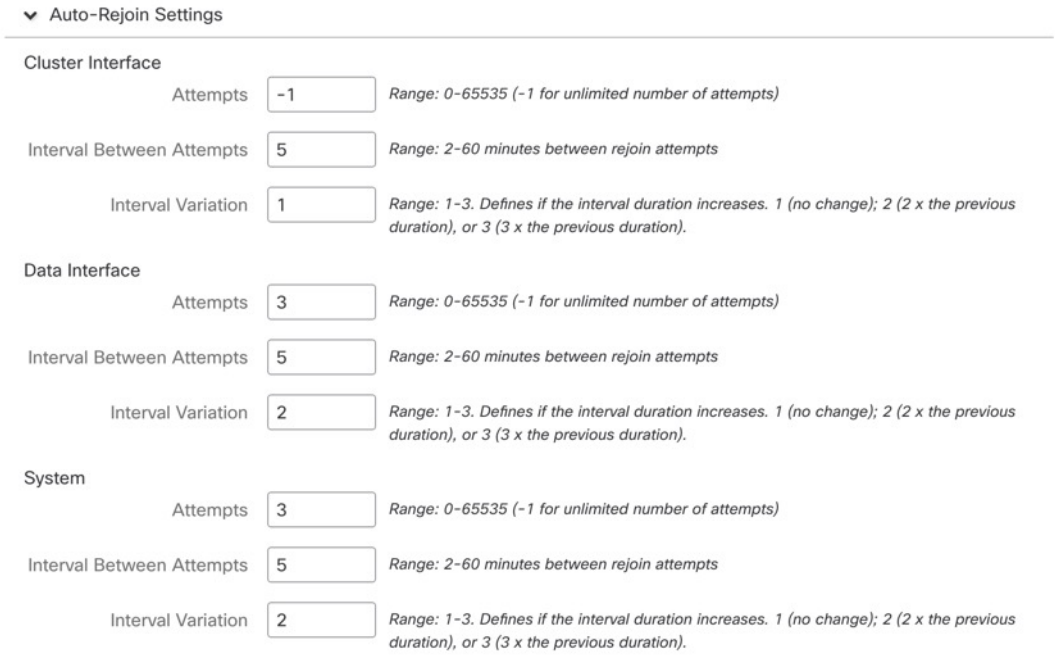
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between 0.3 and 45 seconds. The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the specified number of milliseconds before marking the interface as failed and removing the node from the cluster. If an EtherChannel transitions from a down state to an up state, such as after the switch reloads or enables an EtherChannel, a longer debounce time can prevent the interface from appearing failed on a cluster node if another cluster node bundles the ports more quickly.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 198: Configure Auto-Rejoin settings

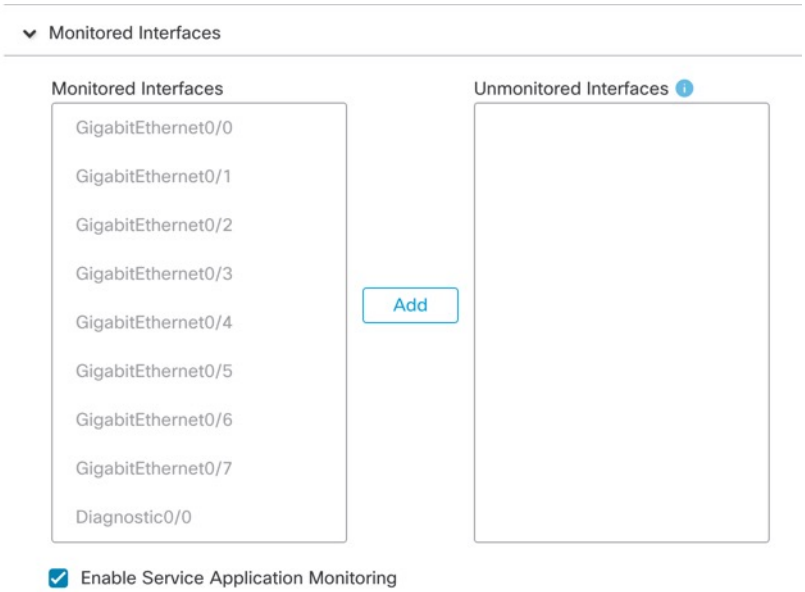


Set values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 199: Configure monitored interfaces



The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The time required to remove a member from the cluster depends on the type of interface and whether the node is established or is joining the cluster. By default, health check is enabled for all interfaces and the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Deploy configuration changes.

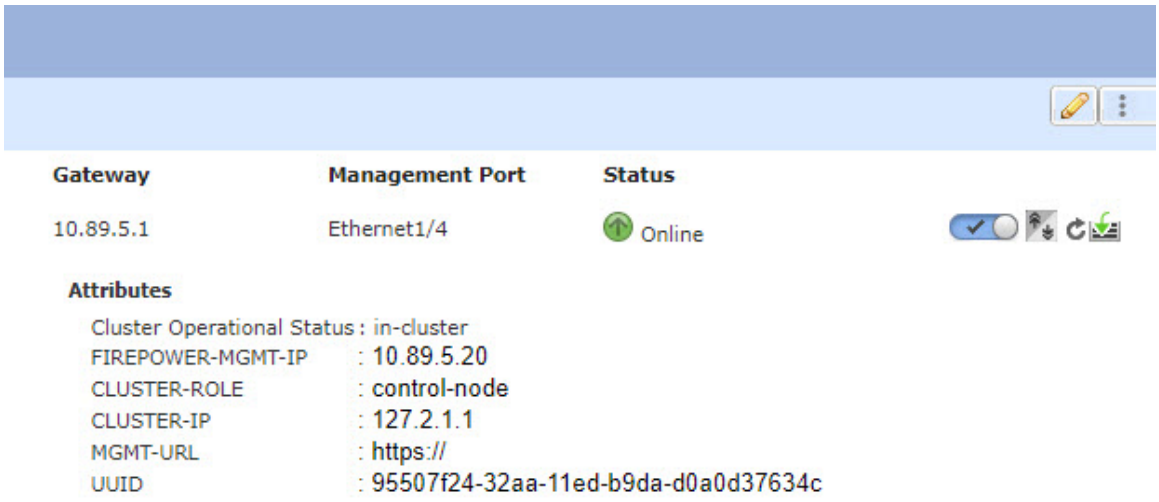
FXOS: Remove a cluster node

The following sections describe how to remove nodes temporarily or permanently from the cluster.

Temporary removal

A cluster node will be automatically removed from the cluster due to a hardware or network failure, for example. This removal is temporary until the conditions are rectified, and it can rejoin the cluster. You can also manually disable clustering.

To check whether a device is currently in the cluster, check the cluster status



For Firewall Threat Defense using the Cloud-Delivered Firewall Management Center, you should leave the device in the Cloud-Delivered Firewall Management Center device list so that it can resume full functionality after you reenable clustering.

- Disable clustering in the application—You can disable clustering using the application CLI. Enter the **cluster remove unit** *name* command to remove any node other than the one you are logged into. The bootstrap configuration remains intact, as well as the last configuration synced from the control node, so you can later re-add the node without losing your configuration. If you enter this command on a data node to remove the control node, a new control node is elected.

  When a device becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, re-enable clustering. The Management interface remains up using the IP address the node received from the bootstrap configuration. However if you reload, and the node is still inactive in the cluster , the Management interface is disabled.

  To reenable clustering, on the Firewall Threat Defense enter **cluster enable**.

- Disable the application instance—
- Shut down the security module/engine—
- Shut down the chassis—

Permanent removal

You can permanently remove a cluster node using the following methods.

For Firewall Threat Defense using the Cloud-Delivered Firewall Management Center, be sure to remove the node from the Cloud-Delivered Firewall Management Center device list after you disable clustering on the chassis.

- Delete the logical device—
- Remove the chassis or security module from service—If you remove a device from service, you can add replacement hardware as a new node of the cluster.

Cloud-Delivered Firewall Management Center: Manage cluster members

After you deploy the cluster, you can change the configuration and manage cluster members.

Add a new cluster member

When you add a new cluster member in FXOS, the Cloud-Delivered Firewall Management Center adds the member automatically.

Before you begin

- Make sure the interface configuration is the same on the replacement unit as for the other chassis.

Follow these steps to add a new cluster member:

Procedure

Step 1 Add the new unit to the cluster in FXOS. Refer to the FXOS configuration guide.

Wait for the new unit to be added to the cluster. Refer to the Firewall Chassis Manager **Logical Devices** screen or use the Firewall Threat Defense **show cluster info** command to view cluster status.

Step 2 The new cluster member is added automatically. To monitor the registration of the replacement unit, view the following:

- **Cluster Status** dialog box (which is available from the **Devices** > **Device Management**, **More** (⋮) icon or from the **Devices** > **Device Management** and choose the **Add**, **Cluster** tab **General** area **Cluster Live Status** link)—A unit that is joining the cluster on the chassis shows "Joining cluster..." After it joins the cluster, the Cloud-Delivered Firewall Management Center attempts to register it, and the status changes to "Available for Registration". After it completes registration, the status changes to "In Sync." If the registration fails, the unit will stay at "Available for Registration". In this case, force a re-registration by clicking **Reconcile**.
- System status **Tasks** —The Cloud-Delivered Firewall Management Center shows all registration events and failures.
- **Devices** > **Device Management** and then choose **Add Device**—When you expand the cluster on the devices listing page, you can see when a unit is registering when it has the loading icon to the left.

Replace a cluster member

You can replace a cluster member in an existing cluster. The Cloud-Delivered Firewall Management Center auto-detects the replacement unit. However, you must manually delete the old cluster member in the Cloud-Delivered Firewall Management Center. This procedure also applies to a unit that was reinitialized; in this case, although the hardware remains the same, it appears to be a new member.

Before you begin

- Make sure the interface configuration is the same on the replacement unit as for other chassis.

Follow these steps to replace a cluster member:

Procedure

Step 1 For a new chassis, if possible, backup and restore the configuration from the old chassis in FXOS.

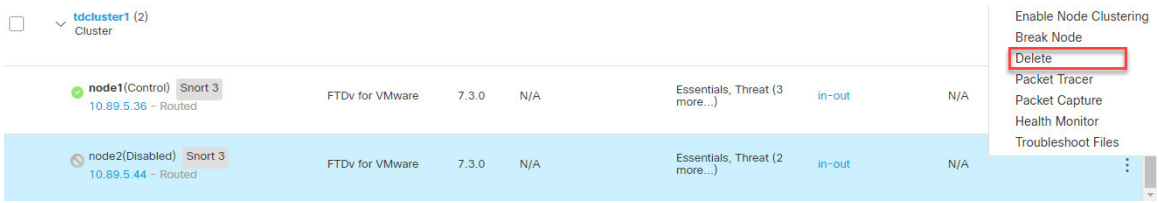
If you are replacing a module in a Firepower 9300, you do not need to perform these steps.

If you do not have a backup FXOS configuration from the old chassis, first perform the steps in Add a new cluster member, on page 964.

For information about all of the below steps, refer to the FXOS configuration guide.

a) Use the configuration export feature to export an XML file containing logical device and platform configuration settings for your Firepower 4100/9300 chassis.
b) Import the configuration file to the replacement chassis.
c) Accept the license agreement.
d) If necessary, upgrade the logical device application instance version to match the rest of the cluster.

Step 2 In the Cloud-Delivered Firewall Management Center for the old unit, choose **Devices** > **Device Management**, **More** (⋮) > **Delete** .



Step 3 Confirm that you want to delete the unit.

The unit is removed from the cluster and from the Cloud-Delivered Firewall Management Center devices list.

Step 4 The new or reinitialized cluster member is added automatically. To monitor the registration of the replacement unit, view the following:

- **Cluster Status** dialog box (**Devices** > **Device Management**, **More** (⋮) icon or **Devices** > **Device Management** and then choose **Add**, **Cluster** page, **General** area, **Cluster Live Status** link)—A unit that is joining the cluster on the chassis shows "Joining cluster..." After it joins the cluster, the Cloud-Delivered Firewall Management Center attempts to register it, and the status changes to "Available for Registration". After it completes registration, the status changes to "In Sync." If the registration fails, the unit will stay at "Available for Registration". In this case, force a re-registration by clicking **Reconcile All**.

- Click **System** (⚙), and then click **Tasks**—The Cloud-Delivered Firewall Management Center shows all registration events and failures.

- **Devices** > **Device Management**, and then click **Add** > **Device**—When you expand the cluster on the devices listing page, you can see when a unit is registering when it has the loading icon to the left.

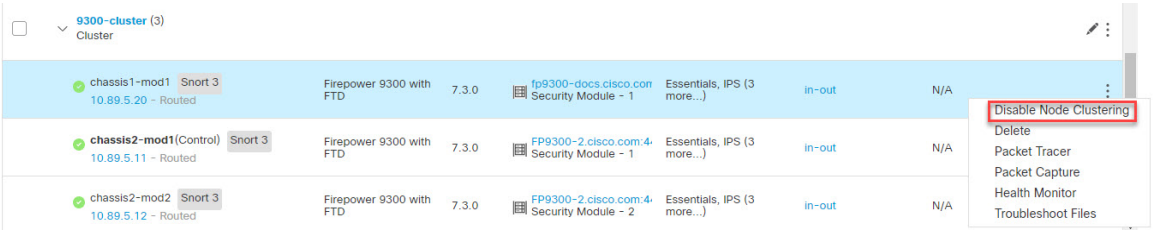
Deactivate a member

You may want to deactivate a member in preparation for deleting the unit, or temporarily for maintenance. This procedure is meant to temporarily deactivate a member; the unit will still appear in the Cloud-Delivered Firewall Management Center device list.

Note When a unit becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, reenable clustering. The Management interface remains up using the IP address the unit received from the bootstrap configuration. However if you reload, and the unit is still inactive in the cluster, the management interface is disabled. You must use the console for any further configuration.

Procedure

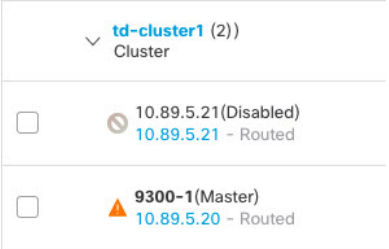
Step 1 For the unit you want to deactivate, choose **Devices** > **Device Management**, **More** (⋮), and then **Disable Clustering**.



You can also deactivate a unit from the **Cluster Status** dialog box (**Devices** > **Device Management**, **More** (⋮), and **Cluster Live Status**).

Step 2 Confirm that you want to disable clustering on the unit.

The unit will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.



Step 3 To re-enable clustering, refer to Rejoin the cluster, on page 967.

Rejoin the cluster

If a unit was removed from the cluster, for example for a failed interface or if you manually disabled clustering, you must manually rejoin the cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management**, **More** (⋮) > **Enable Clustering**.



You can also reactivate a unit from the **Cluster Status** dialog box (**Devices** > **Device Management**, **More** (⋮) > **Cluster Live Status**).

Step 2 Confirm that you want to enable clustering on the unit.

Unregister a data node

If you need to permanently remove a cluster node (for example, if you remove a module on the Firepower 9300, or remove a chassis), then you should unregister it from the Cloud-Delivered Firewall Management Center.

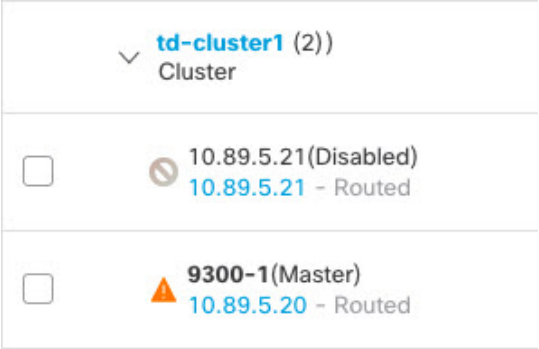
Do not unregister the node if it is still a healthy part of the cluster, or if you only want to disable the node temporarily. To remove it permanently from the cluster in FXOS, see FXOS: Remove a cluster node, on page 962. If you unregister it from the Cloud-Delivered Firewall Management Center, and it is still part of the cluster, it will continue to pass traffic, and could even become the control node—a control node that the Cloud-Delivered Firewall Management Center can no longer manage.

Before you begin

To manually deactivate the node, see Deactivate a member, on page 966. Before you unregister a node, the node must be inactive, either manually or because of a health failure.

Procedure

Step 1 Make sure the node is ready to be unregistered from the Cloud-Delivered Firewall Management Center. On **Devices** > **Device Management**, make sure the node shows **(Disabled)**.



You can also view each node's status on the **Cluster Status** dialog box available from **More** (⋮). If the status is stale, click **Reconcile All** on the **Cluster Status** dialog box to force an update.

Step 2 In the Cloud-Delivered Firewall Management Center for the data node you want to delete, choose **Devices** > **Device Management** and then click **Add** > **DeviceMore** (⋮) > **Unregister**.

Step 3 Confirm that you want to unregister the node.

The node is removed from the cluster and from the Cloud-Delivered Firewall Management Center devices list.

Change the control unit

The best method to change the control unit is to disable clustering on the control unit, wait for a new control election, and then re-enable clustering.

Caution If you must specify the *exact* unit you want to become the control unit, use the procedure in this section. Note that for centralized features, if you force a control unit change, then all connections are dropped, and you have to re-establish the connections on the new control unit.

Procedure

Step 1 Open the **Cluster Status** dialog box by choosing **Devices** > **Device Management**, **More** (⋮) > **Cluster Live Status**.

You can also access the **Cluster Status** dialog box from **Devices** > **Device Management** and choose the **Add**, **Cluster** page, **General** area, and **Cluster Live Status** link.

Step 2 For the unit you want to become the control unit, choose **More** (⋮) > **Change Role to Control**.

Step 3 You are prompted to confirm the role change. Check the check box, and click **OK**.

Reconcile cluster members

If a cluster member fails to register, you can reconcile the cluster membership from the chassis to the Cloud-Delivered Firewall Management Center. For example, a data unit might fail to register if the Cloud-Delivered Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

Step 1 Choose **Devices** > **Device Management**, **More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

You can also open the **Cluster Status** dialog box from the **Devices** > **Device Management** page **General** area, **Cluster Live Status** link.

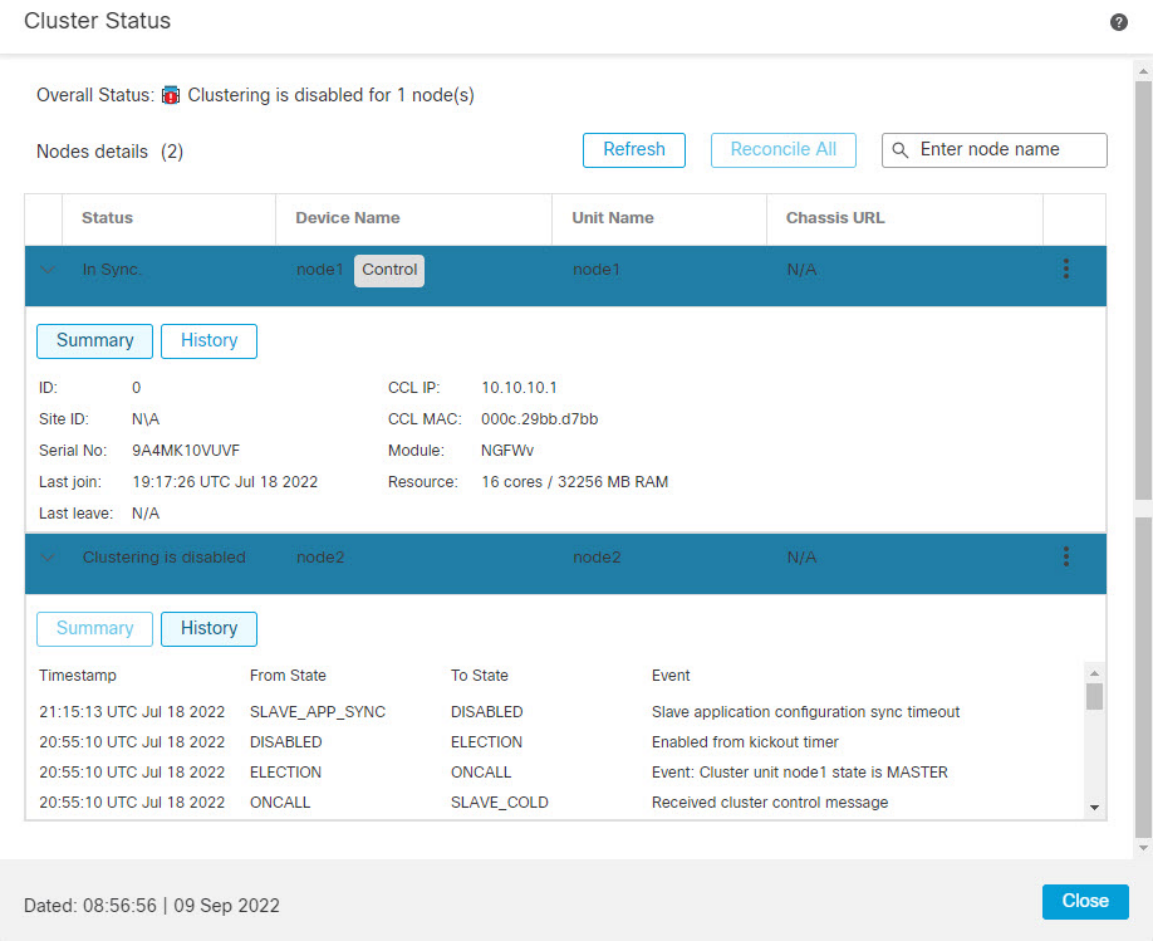
Step 2 Click **Reconcile All**.

For more information about the cluster status, refer to Cloud-Delivered Firewall Management Center: Monitor the cluster, on page 969.

Cloud-Delivered Firewall Management Center: Monitor the cluster

You can monitor the cluster in Cloud-Delivered Firewall Management Center and at the Firewall Threat Defense CLI.

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management**, **More** (⋮) icon or from the **Devices** > **Device Management** and then chose **Add**, **Cluster** page, **General** area, and **Cluster Live Status** link.



The Control unit has a graphic indicator identifying its role.

Cluster member **Status** includes the following states:

- In Sync.—The unit is registered with the Cloud-Delivered Firewall Management Center.
- Pending Registration—The unit is part of the cluster, but has not yet registered with the Cloud-Delivered Firewall Management Center. If a unit fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The unit is registered with the Cloud-Delivered Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the unit from the cluster.
- Joining cluster...—The unit is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Cloud-Delivered Firewall Management Center.

For each unit, you can view the **Summary** or the **History**.

For each unit from the **More** (⋮) menu , you can perform the following status changes:

- **Disable Clustering**
- **Enable Clustering**

- **Change Role to Control**

- Click **System** (⚙), and then click **Tasks**.

  The **Tasks** page shows updates of the Cluster Registration task as each unit registers.

- **Devices** > **Device Management** and then choose **Add Device** *cluster_name*.

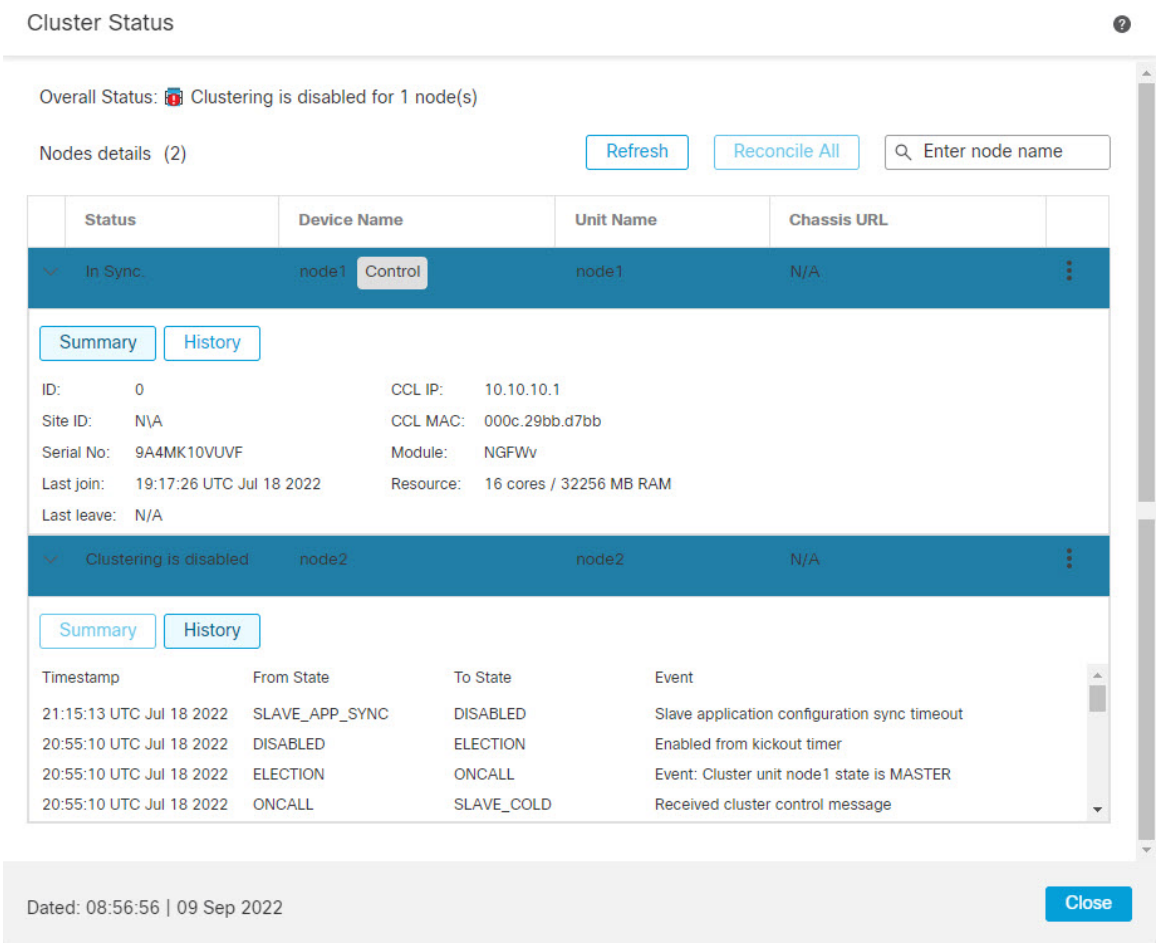
  When you expand the cluster on the devices listing page, you can see all member units, including the control unit shown with its role next to the IP address. For units that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.



The Control unit has a graphic indicator identifying its role.

Cluster member **Status** includes these states:

- In Sync.—The unit is registered with the Cloud-Delivered Firewall Management Center.
- Pending Registration—The unit is part of the cluster, but has not yet registered with the Cloud-Delivered Firewall Management Center. If a unit fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The unit is registered with the Cloud-Delivered Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the unit from the cluster.
- Joining cluster...—The unit is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Cloud-Delivered Firewall Management Center.

For each unit, you can view the **Summary** or the **History**.

For each unit from the **More** (⋮) menu, you can perform these status changes:

- **Disable Clustering**
- **Enable Clustering**
- **Change Role to Control**

Tasks page

Click **System** (⚙), and then click **Tasks**.

The **Tasks** page shows updates of the Cluster Registration task as each unit registers.

Device management page

Devices > **Device Management** and then choose **Add Device** *cluster_name*.

When you expand the cluster on the devices listing page, you can see all member units, including the control unit shown with its role next to the IP address. For units that are still registering, you can see the loading icon.

CLI commands for cluster monitoring

show cluster {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

To view aggregated data for the entire cluster or other information, use the **show cluster** command.

show cluster info [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

To view cluster information, use the **show cluster info** command.

Cluster Health Monitor Dashboard

When a Firewall Threat Defense is the control node of a cluster, the Cloud-Delivered Firewall Management Center collects various metrics periodically from the device metric data collector.

The cluster health monitor is comprised of these components:

- Overview dashboard: Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Cloud-Delivered Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.



Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard: Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.

- Member Performance dashboard: Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard: Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links: Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range: An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard: Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

Viewing cluster health

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

Ensure you have created a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Follow these steps to view the cluster health monitor:

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed cluster devices.

Step 3 Click on the cluster name to view the cluster health statistics.

The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution: Traffic and packet distribution across the cluster nodes.
- Member Performance: Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL: Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 View device monitor health and performance metrics in several predefined dashboards.

The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.

- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics that are related to the Snort process.
- ASP drops: Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. Refer to Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 63: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number
Packets	Packet distribution count in the cluster for every second.	number

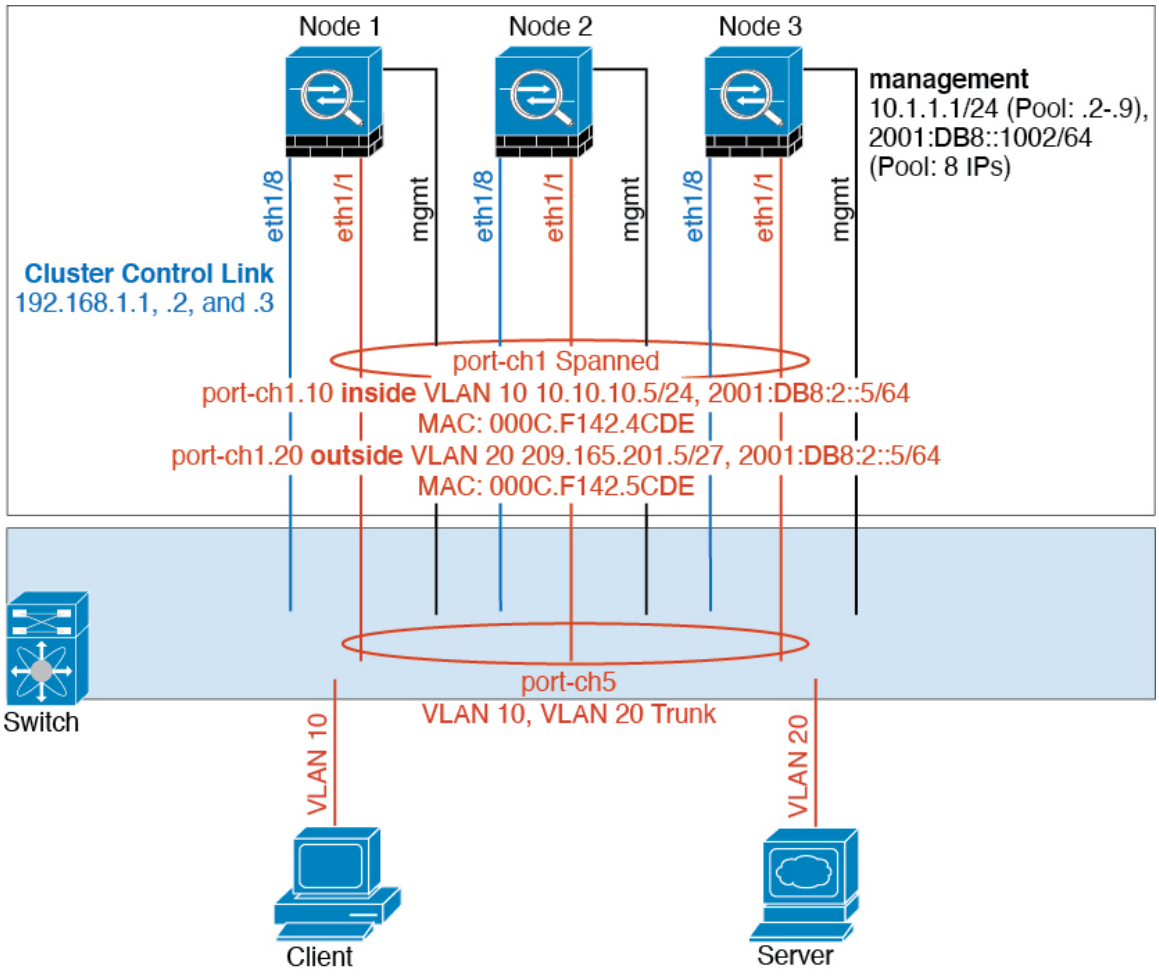
Examples for Clustering

These examples include typical deployments.

Firewall on a stick

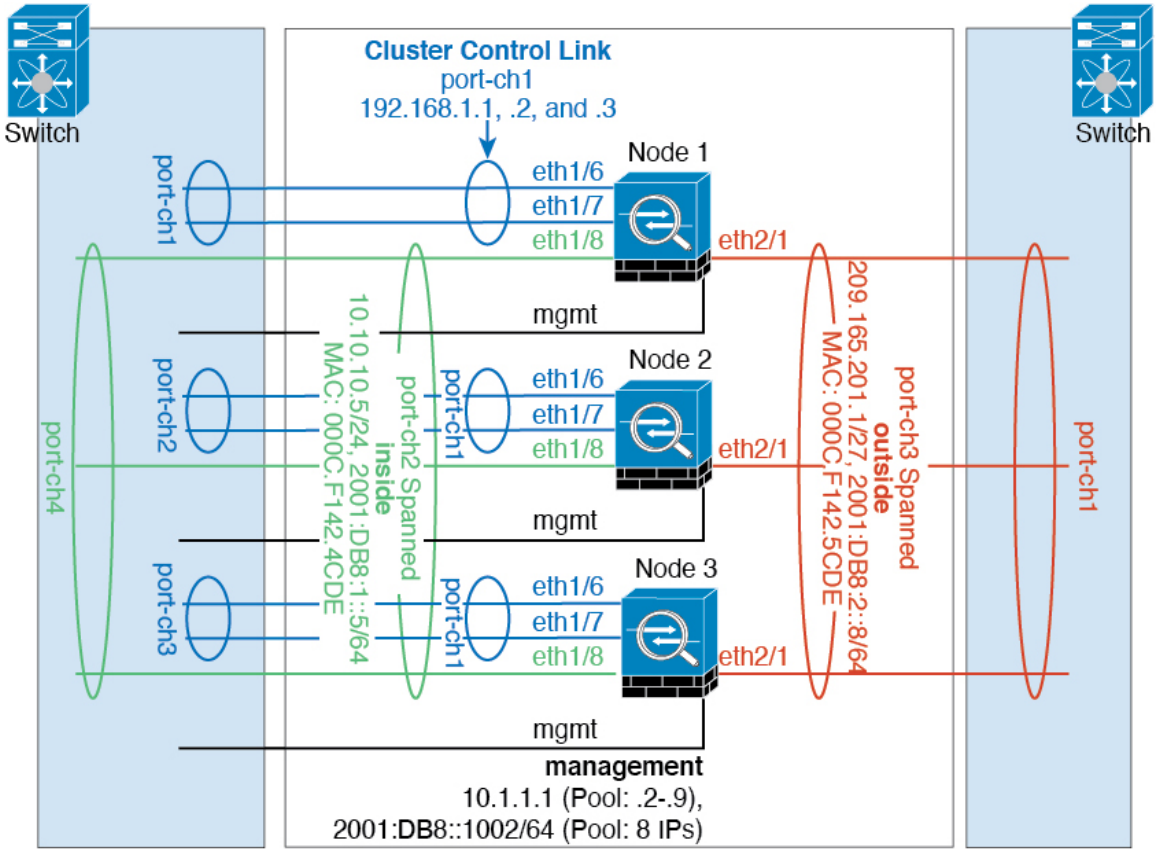
Data traffic from different security domains are associated with different VLANs. For example, VLAN 10 for the inside network and VLAN 20 for the outside network. Each Threat Defense has a single physical port connected to the external switch or router. Trunking is enabled so that all packets on the physical link are 802.1q encapsulated. The Threat Defense is the firewall between VLAN 10 and VLAN 20.

When using Spanned EtherChannels, all data links are grouped into one EtherChannel on the switch side. If the Threat Defense becomes unavailable, the switch will rebalance traffic between the remaining units.



Traffic segregation

Provides the necessary bootstrap and interface configurations to implement traffic segregation using spanned EtherChannels in a cluster environment.



As shown in the diagram above, there is one Spanned EtherChannel on the left side that connects to the inside switch, and the other on the right side to outside switch. You can also create VLAN subinterfaces on each EtherChannel if desired.

Reference for clustering

This section includes more information about how clustering operates.

Firewall Threat Defense features and clustering

Some Firewall Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported features with clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Cloud-Delivered Firewall Management Center GUI. See FlexConfig policies, on page 1309.

- Remote access VPN (SSL VPN and IPsec VPN)
- DHCP client, server, and proxy. DHCP relay is supported.
- Virtual Tunnel Interfaces (VTIs)
- High Availability
- Integrated Routing and Bridging
- Cloud-Delivered Firewall Management Center UCAPL/CC mode
- DHCP client, server, and proxy. DHCP relay is supported.

Centralized features for clustering

The following features are only supported on the control unit, and are not scaled for the cluster.



Note Traffic for centralized features is forwarded from member units to the control unit over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control units before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control unit.

For centralized features, if the control unit fails, all connections are dropped, and you have to re-establish the connections on the new control unit.

- Site-to-site VPN
- The following application inspections:
  - DCERPC
  - NetBIOS
  - RSH
  - SUNRPC
  - TFTP
  - XDMCP
- Dynamic routing
- Static route monitoring

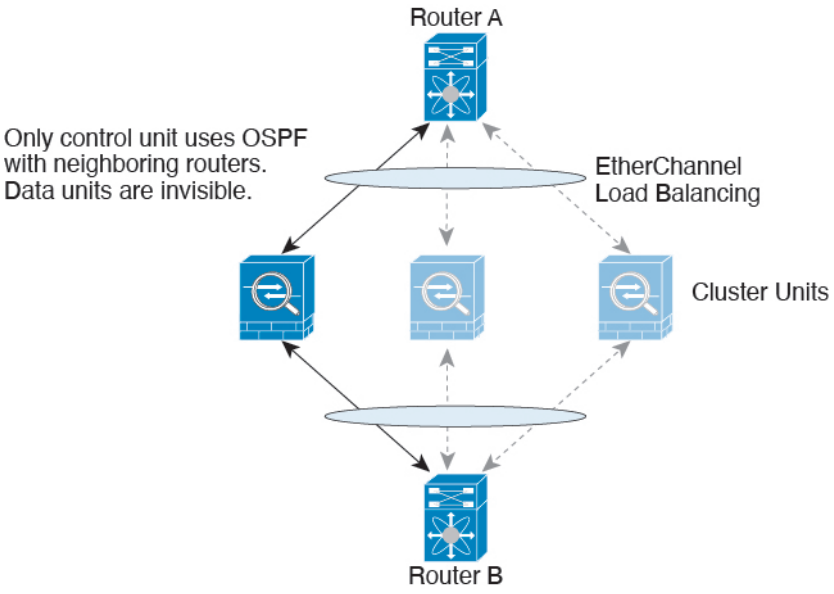
Connection settings

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

Dynamic routing and clustering

The routing process only runs on the control unit, and routes are learned through the control unit and replicated to secondaries. If a routing packet arrives at a data unit, it is redirected to the control unit.

Figure 200: Dynamic Routing



After the data units learn the routes from the control unit, each unit makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the control unit to data units. If there is a control unit switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

FTP and clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

Multicast routing and clustering

The control unit handles all multicast routing packets and data packets until fast-path forwarding is established. After the connection is established, each data unit can forward multicast data packets.

NAT and clustering

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different Firewall Threat Defenses in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When

a packet arrives at the Firewall Threat Defense that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- PAT with Port Block Allocation—See the following guidelines for this feature:
  - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
  - Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
  - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
  - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.
- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
  - FTP

- RSH
- SQLNET
- TFTP
- XDMCP
- SIP

- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

SIP inspection and clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and clustering

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

TLS/SSL connections and clustering

The decryption states of TLS/SSL connections are not synchronized, and if the connection owner fails, then the decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Cisco TrustSec and clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

VPN and clustering

Site-to-site VPN is a centralized feature; only the control unit supports VPN connections.

Note Remote access VPN is not supported with clustering.

VPN functionality is limited to the control unit and does not take advantage of the cluster high availability capabilities. If the control unit fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control unit is elected, you must reestablish the VPN connections.

When you connect a VPN tunnel to a Spanned interface address, connections are automatically forwarded to the control unit.

VPN-related keys and certificates are replicated to all units.

Performance scaling factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

For example, for TCP throughput, the Firepower 9300 with 3 SM-40 modules can handle approximately 135 Gbps of real world firewall traffic when running alone. For 2 chassis, the maximum combined throughput will be approximately 80% of 270 Gbps (2 chassis x 135 Gbps): 216 Gbps.

Control unit election

The control unit election is a automated process where cluster members determine which unit manages centralized features.

Election process and rules

Members of the cluster communicate over the cluster control link to elect a control unit as follows:

1. When you deploy the cluster, each unit broadcasts an election request every 3 seconds.
2. Any other units with a higher priority respond to the election request; the priority is set when you deploy the cluster and is not configurable.
3. If after 45 seconds, a unit does not receive a response from another unit with a higher priority, then it becomes the control unit.
4. If a unit later joins the cluster with a higher priority, it does not automatically become the control unit; the existing control unit always remains as the control unit unless it stops responding, at which point a new control unit is elected.
5. In a *split brain* scenario when there are temporarily multiple control units, then the unit with highest priority retains the role while the other units return to data unit roles.

Note You can manually force a unit to become the control unit. For centralized features, if you force a control unit change, then all connections are dropped, and you have to re-establish the connections on the new control unit.

High availability within the cluster

Clustering provides high availability by monitoring chassis, unit, and interface health and by replicating connection states between units.

Chassis-application monitoring

Chassis-application health monitoring is always enabled. The Firepower 4100/9300 chassis supervisor checks the Firewall Threat Defense application periodically (every second). If the Firewall Threat Defense device is up and cannot communicate with the Firepower 4100/9300 chassis supervisor for 3 seconds, the Firewall Threat Defense device generates a syslog message and leaves the cluster.

If the Firepower 4100/9300 chassis supervisor cannot communicate with the application after 45 seconds, it reloads the Firewall Threat Defense device. If the Firewall Threat Defense device cannot communicate with the supervisor, it removes itself from the cluster.

Unit health monitoring

Each unit periodically sends a broadcast keepaliveheartbeat packet over the cluster control link. If the control node does not receive any keepaliveheartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining node.

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role. See Control unit election, on page 982 for more information.

Interface monitoring

Each node monitors the link status of all hardware interfaces in use, and reports status changes to the control node. For clustering on multiple chassis, Spanned EtherChannels use the cluster Link Aggregation Control Protocol (cLACP). Each chassis monitors the link status and the cLACP protocol messages to determine if the port is still active in the EtherChannel, and informs the Firewall Threat Defense application if the interface is down. When you enable health monitoring, all physical interfaces are monitored by default (including the main EtherChannel for EtherChannel interfaces). Only named interfaces that are in an Up state can be monitored. For example, all member ports of an EtherChannel must fail before a *named* EtherChannel is removed from the cluster. You can optionally disable monitoring per interface.

If a monitored interface fails on a particular node, but it is active on other nodes, then the node is removed from the cluster. The amount of time before the Firewall Threat Defense device removes a node from the cluster depends on whether the node is an established member or is joining the cluster. The Firewall Threat Defense device does not monitor interfaces for the first 90 seconds that a node joins the cluster. Interface status changes during this time will not cause the Firewall Threat Defense device to be removed from the cluster. For an established member, the node is removed after 500 ms.

For clustering on multiple chassis, if you add or delete an EtherChannel from the cluster, interface health-monitoring is suspended for 95 seconds to ensure that you have time to make the changes on each chassis.

Decorator application monitoring

When you install a decorator application on an interface, such as the Radware DefensePro application, then both the Firewall Threat Defense device and the decorator application must be operational to remain in the cluster. The unit does not join the cluster until both applications are operational. Once in the cluster, the unit monitors the decorator application health every 3 seconds. If the decorator application is down, the unit is removed from the cluster.

Status after failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The Firewall Threat Defense automatically tries to rejoin the cluster, depending on the failure event.



Note When the Firewall Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoining the cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

Rejoining scenarios and automatic retry behavior

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firewall Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firewall Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firewall Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from Cloud-Delivered Firewall Management Center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members

are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.

- Failed Chassis-Application Communication—When the Firewall Threat Defense application detects that the chassis-application health has recovered, it tries to rejoin the cluster automatically.

Data path connection state replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 64: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the cluster manages connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.

- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.

- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.



Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

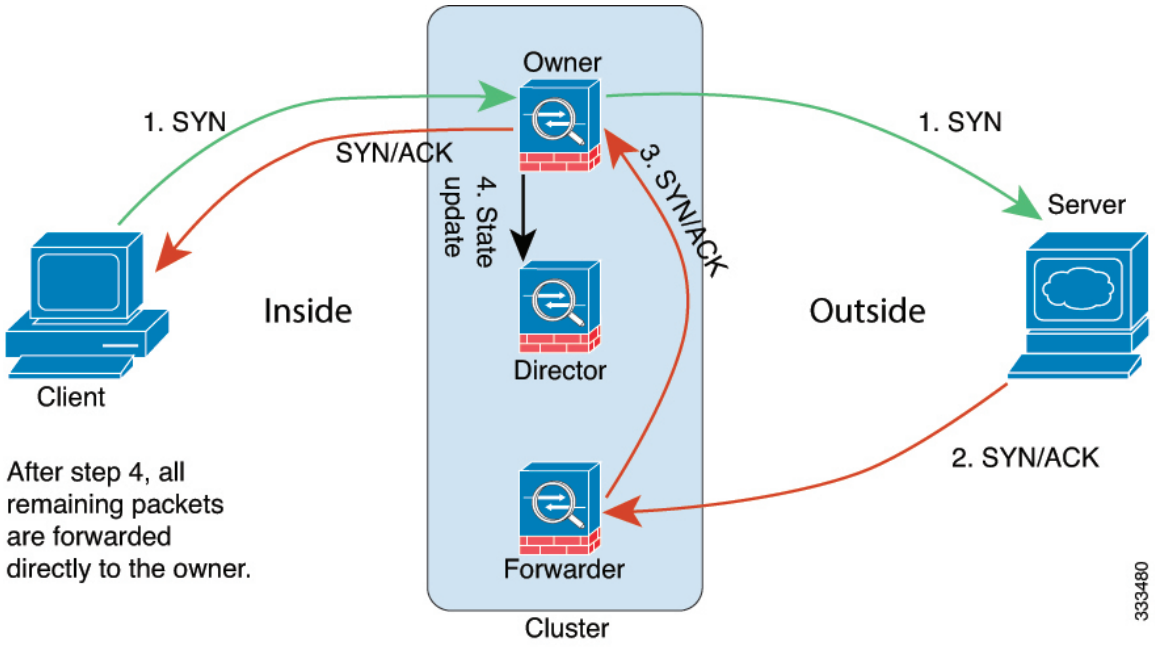
Port Address Translation Connections

New connection ownership

When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. If a reverse flow arrives at a different node, it is redirected back to the original node.

Sample data flow for TCP

The following example shows the establishment of a new connection.

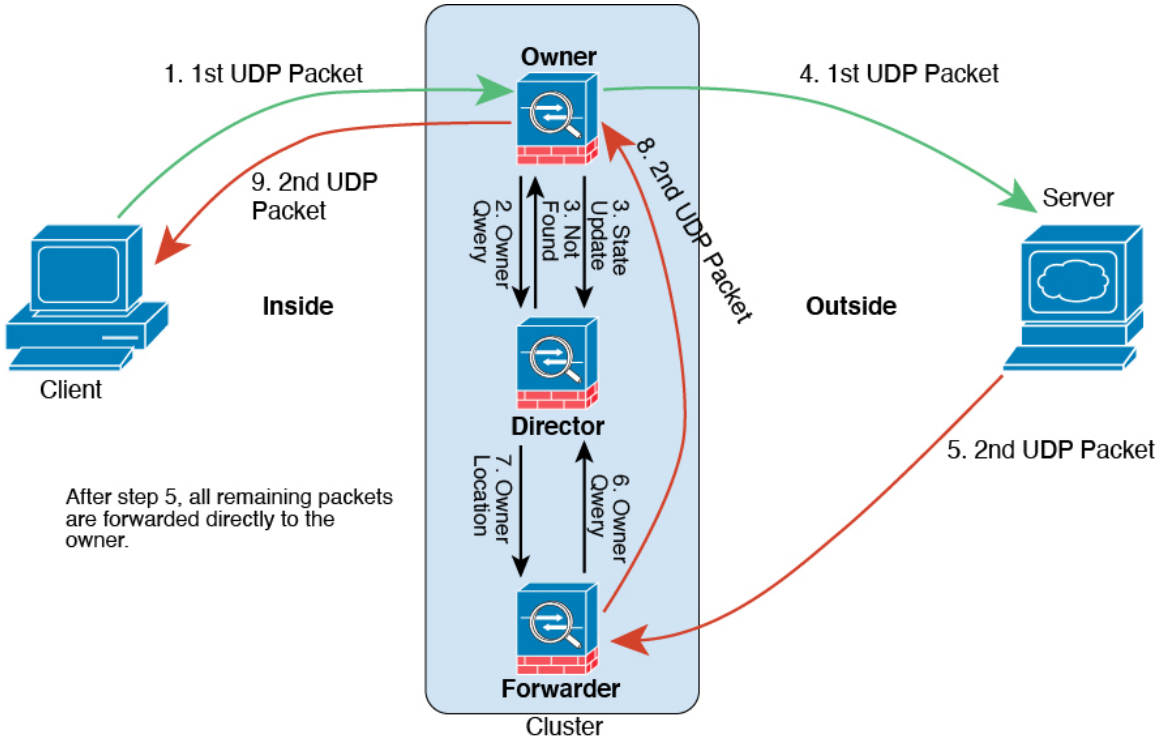


1. The SYN packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different Firewall Threat Defense (based on the load balancing method). This Firewall Threat Defense is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample data flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. ***Figure 201: ICMP and UDP Data Flow***



   The first UDP packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for clustering

Table 65:

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
MTU ping test from data node on node join	20241030	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size twice the cluster control link MTU. Previously, only the control node sent a ping. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again. Added/modified commands: **show cluster history**. See: Clustering for the Secure Firewall 3100/4200, Clustering for Threat Defense Virtual in a Private Cloud, Clustering for Threat Defense Virtual in a Public Cloud, Clustering for the Firepower 4100/9300
Cluster health monitor settings.	20221213	Any	You can now edit cluster health monitor settings. New/Modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard.	20221213	Any	You can now view cluster health on the cluster health monitor dashboard. New/Modified screens: **System** > **Health** > **Monitor**
Support for 16-node clusters.	20220609	7.2.0	You can now configure 16 node clusters for the Firepower 4100/9300. Previously, the maximum was 6 units. New/Modified screens: none. Supported platforms: Firepower 4100/9300
Cluster deployment for firewall changes completes faster.	20220609	7.2.0	Cluster deployment for firewall changes now completes faster. New/Modified screens: none.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Improved PAT port block allocation for clustering.	20220609	7.0.3	The improved PAT port block allocation ensures that the control unit keeps ports in reserve for joining nodes, and proactively reclaims unused ports. To best optimize the allocation, you can set the maximum nodes you plan to have in the cluster using the **cluster-member-limit** command using FlexConfig. The control unit can then allocate port blocks to the planned number of nodes, and it will not have to reserve ports for extra nodes you don't plan to use. The default is 16 nodes. You can also monitor syslog 747046 to ensure that there are enough ports available for a new node. New/Modified commands: **cluster-member-limit** (FlexConfig), **show nat pool cluster** [**summary**], **show nat pool ip detail**
Cluster deployment for Snort changes completes faster, and fails faster when there is an event.	20220609	7.0.3	Cluster deployment for Snort changes now completes faster. Also, when a cluster has an event that causes a Cloud-Delivered Firewall Management Center deployment to fail, the failure now occurs more quickly. New/Modified screens: none.
Improved cluster management.	20220609	7.0.3	Cloud-Delivered Firewall Management Center has improved cluster management functionality that formerly you could only accomplish using the CLI, including: • Enable and disable cluster units • Show cluster status from the Device Management page, including History and Summary per unit • Change the role to the control unit New/Modified screens: • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** > **General** area **> Cluster Live Status** link **Cluster Status** Supported platforms: Firepower 4100/9300

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Multi-instance clustering.	20220609	7.0.3	You can now create a cluster using container instances. On the Firepower 9300, you must include one container instance on each module in the cluster. You cannot add more than one container instance to the cluster per security engine/module. We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. New/Modified FXOS commands: **set port-type cluster** New/modified Firepower Chassis Manager screens: • **Logical Devices** > **Add Cluster** • **Interfaces** > **All Interfaces** > **Add New** drop-down menu**Subinterface Type field** Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Configuration sync to data units in parallel.	20220609	7.0.3	The control unit now syncs configuration changes with data units in parallel by default. Formerly, synching occurred sequentially. New/Modified screens: none.
Messages for cluster join failure or eviction added to **show cluster history**.	20220609	7.0.3	New messages were added to the **show cluster history** command for when a cluster unit either fails to join the cluster or leaves the cluster. New/Modified commands: **show cluster history** New/Modified screens: none.
Initiator and responder information for Dead Connection Detection (DCD), and DCD support in a cluster.	20220609	7.0.3	If you enable Dead Connection Detection (DCD), you can use the **show conn detail** command to get information about the initiator and responder. Dead Connection Detection allows you to maintain an inactive connection, and the **show conn** output tells you how often the endpoints have been probed. In addition, DCD is now supported in a cluster. New/Modified commands: **show conn** (output only). Supported platforms: Firewall Threat Defense on the Firepower 4100/9300

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Adding clusters is easier.	20220609	7.0.3	You can now add any unit of a cluster to the Cloud-Delivered Firewall Management Center, and the other cluster units are detected automatically. Formerly, you had to add each cluster unit as a separate device, and then group them into a cluster. Adding a cluster unit is also now automatic. Note that you must delete a unit manually. New/Modified screens: **Devices** > **Device Management** > **Add** drop-down menu **Device** > **Add Device** dialog box **Devices** > **Device Management** > **Cluster** tab**General** area**Cluster Registration Status** **Current Cluster Summary** link **Cluster Status** dialog box Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Support for site-to-site VPN with clustering as a centralized feature.	20220609	7.0.3	You can now configure site-to-site VPN with clustering. Site-to-site VPN is a centralized feature; only the control unit supports VPN connections. Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Automatically rejoin the cluster after an internal failure.	20220609	7.0.3	Formerly, many internal error conditions caused a cluster unit to be removed from the cluster, and you were required to manually rejoin the cluster after resolving the issue. Now, a unit will attempt to rejoin the cluster automatically at the following intervals: 5 minutes, 10 minutes, and then 20 minutes. Internal failures include: application sync timeout; inconsistent application statuses; and so on. New/Modified command: **show cluster info auto-join** No modified screens. Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Clustering on multiple chassis for 6 modules; Firepower 4100 support.	20220609	7.0.3	With FXOS 2.1.1, you can now enable clustering on multiple chassis of the Firepower 9300 and 4100. For the Firepower 9300, you can include up to 6 modules. For example, you can use 1 module in 6 chassis, or 2 modules in 3 chassis, or any combination that provides a maximum of 6 modules. For the Firepower 4100, you can include up to 6 chassis. **Note** Inter-site clustering is also supported. However, customizations to enhance redundancy and stability, such as site-specific MAC and IP addresses, director localization, site redundancy, and cluster flow mobility, are only configurable using the FlexConfig feature. No modified screens. Supported platforms: Firewall Threat Defense on the Firepower 4100/9300

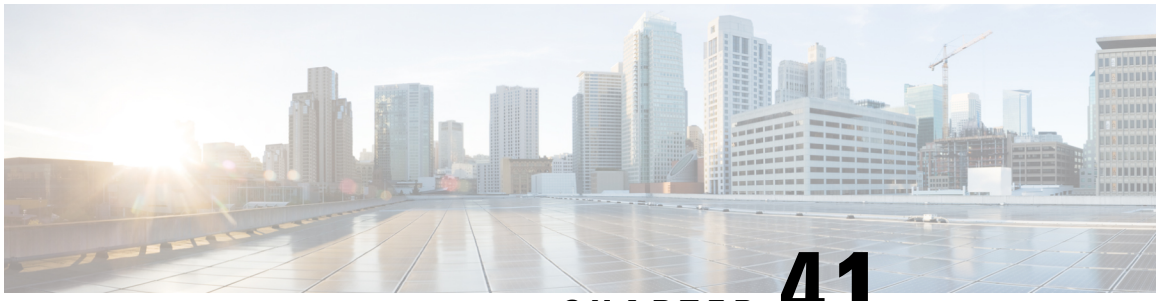
Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Clustering on multiple modules with one Firepower 9300 chassis.	20220609	7.0.3	You can cluster up to 3 security modules within the Firepower 9300 chassis. All modules in the chassis must belong to the cluster. New/Modified screens: **Devices** > **Device Management** > **Add** > **Add Cluster** **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense on the Firepower 9300



PART XII

Interfaces and Device Settings

- Interface Overview, on page 997
- Regular Firewall Interfaces, on page 1039
- Inline Sets and Passive Interfaces, on page 1119
- Device Settings, on page 1133
- DHCP and DDNS, on page 1223
- Platform Settings, on page 1241
- FlexConfig policies, on page 1309



CHAPTER 41

Interface Overview

The Firewall Threat Defense device includes data interfaces that you can configure in different modes, as well as a management interface.

- Management interface, on page 997
- Interface Mode and Types, on page 999
- Security zones and interface groups, on page 999
- Auto-MDI/MDIX feature, on page 1001
- Redundant interfaces (deprecated), on page 1001
- Default settings for interfaces, on page 1001
- Create security zone and interface group objects, on page 1002
- Enable the physical interface and configure Ethernet settings, on page 1003
- Configure EtherChannel interfaces, on page 1006
- Sync Interface Changes with the Cloud-Delivered Firewall Management Center, on page 1014
- Network module management for Secure Firewall 3100/4200, on page 1015
- Merge the management and diagnostic interfaces, on page 1030
- History for interfaces, on page 1035

Management interface

A management interface is a shared physical interface that enables communication between the diagnostic logical interface and the management logical interface on security appliances.

Interface configuration details

In Version 7.3 and earlier, the physical management interface is shared between the Diagnostic logical interface and the Management logical interface. In Version 7.4 and later, the diagnostic interface is merged with management for a simplified user experience.

Management interface

The Management interface is separate from the other interfaces on the device. It is used to set up and register the device to the Cloud-Delivered Firewall Management Center. It uses its own IP address and static routing.

Management interface configuration and monitoring

You can configure its settings at the CLI using the **configure network** command. You can also view its status on the **Devices** > **Device Management** page. If you change the IP address at the CLI after you add it to the Cloud-Delivered Firewall Management Center, you can match the IP address in the Cloud-Delivered Firewall Management Center in the **Devices** > **Device Management** area.

You can alternatively manage the Firewall Threat Defense using a data interface instead of the management interface.

Diagnostic interface

Diagnostic interface shares a physical port with the Management interface but requires different IP addresses on the same network. A diagnostic interface allows only management traffic.

Legacy diagnostic interface support

For new devices using 7.4 and later, you cannot use the legacy Diagnostic interface. Only the merged Management interface is available.

If you upgraded to 7.4 or later, and you did not have any configuration for the Diagnostic interface, then the interfaces will merge automatically.

If you upgraded to 7.4 or later, and you have configuration for the Diagnostic interface, then you have the choice to merge the interfaces manually, or you can continue to use the separate Diagnostic interface. Support for the Diagnostic interface will be removed in a later release, so you should plan to merge the interfaces as soon as possible. To manually merge the Management and Diagnostic interfaces, see Merge the management and diagnostic interfaces, on page 1030. Configurations that prevent an automatic merge include these:

- A data interface named "management"—This name is reserved for use with the merged Management interface.
- IP Address on Diagnostic
- DNS enabled on Diagnostic
- Syslog, SNMP, RADIUS or AD (for remote access VPN) source interface is Diagnostic
- RADIUS or AD (for remote access VPN) with no source interface specified, and there is at least one interface configured as management-only (including Diagnostic)—The default route lookup for these services has changed from the management-only routing table to the data routing table, with no fallback to management. Therefore, you cannot use a management-only interface other than Management.
- Static routes on Diagnostic
- Dynamic routing on Diagnostic
- HTTP server on Diagnostic
- ICMP on Diagnostic
- DDNS for Diagnostic
- FlexConfig using Diagnostic

For more information about how the legacy Diagnostic interface operates, see the 7.3 version of this guide.

Interface Mode and Types

You can deploy Firewall Threat Defense interfaces in two modes: Regular firewall mode and IPS-only mode. You can include both firewall and IPS-only interfaces on the same device.

Regular Firewall Mode

Firewall mode interfaces subject traffic to firewall functions such as maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization. You can also optionally configure IPS functions for this traffic according to your security policy.

The types of firewall interfaces you can configure depends on the firewall mode set for the device: routed or transparent mode. See Transparent or Routed Firewall Mode, on page 153 for more information.

- Routed mode interfaces (routed firewall mode only)—Each interface that you want to route between is on a different subnet.
- Bridge group interfaces (routed and transparent firewall mode)—You can group together multiple interfaces on a network, and the Firewall Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. In routed mode, the Firewall Threat Defense device routes between BVIs and regular routed interfaces. In transparent mode, each bridge group is separate and cannot communicate with each other.

IPS-Only Mode

You can configure your device in either a passive or inline IPS deployment. In a passive deployment, you deploy the system out-of-band from the flow of network traffic. In an inline deployment, you configure the system transparently on a network segment by binding two interfaces together.

Security zones and interface groups

You can assign each interface to a security zone, an interface group, or both. You then apply your security policy based on zones or groups. For example, you can assign the "inside" interface on one or more devices to the "inside" zone, and the "outside" interfaces to the "outside" zone. You can then configure your access control policy to enable traffic to go from the inside zone to the outside zone for every device using the same zones.

Security zones and interface groups functionality

To view the interfaces that belong to each object, choose **Objects** > **Object Management** and click **Interface**. This page lists the security zones and interface groups configured on your managed devices. You can expand each interface object to view the type of interfaces in each interface object.

Note Policies that apply to **any** zone (a global policy) apply to interfaces in zones as well as any interfaces that are not assigned to a zone.

Note The Management interface does not belong to a zone or interface group.

Security zones and interface groups have different membership and usage characteristics:

- Security zones—An interface can belong to only one security zone.
- Interface groups—An interface can belong to multiple interface groups (and to one security zone).

  You can use interface groups in NAT policies, prefilter policies, and QoS policies, as well as features that let you specify the interface name directly, such as Syslog servers or DNS servers.

Some policies only support security zones, while other policies support zones and groups. Unless you need the functionality an interface group provides, you should default to using security zones because security zones are supported for all features.

You cannot change an existing security zone to an interface group or vice-versa. Instead, you must create a new interface object.

Note Although tunnel zones are not interface objects, you can use them in place of security zones in certain configurations; see Using tunnel zones to apply access control at the tunnel level, on page 2534.

Interface objects are categorized by these types:

- Passive—For IPS-only passive or ERSPAN interfaces.
- Inline—For IPS-only inline set interfaces.
- Switched—For regular firewall bridge group interfaces.
- Routed—For regular firewall routed interfaces.
- ASA—(Security zones only) For legacy ASA FirePOWER device interfaces.
- Management—(Interface groups only) For management-only interfaces.
- Loopback—(Interface groups only) For loopback interfaces.

All interfaces in an interface object must be of the same type. After you create an interface object, you cannot change the type of interfaces it contains.

The interface (or zone name) itself does not provide any default behavior in regarding the security policy. We recommend using names that are self-describing to avoid mistakes in future configuration. A good name signifies a logical segment or traffic specification, for example:

- Names of internal interfaces—InsideV110, InsideV160, InsideV195
- Names of DMZ interfaces—DMZV11, DMZV12, DMZV-TEST
- Names of external interfaces—Outside-ASN78, Outside-ASN91

Auto-MDI/MDIX feature

Auto-MDI/MDIX is a feature that eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. The feature is enabled by default with auto-negotiation on RJ-45 interfaces and requires either speed or duplex to be set to auto-negotiate. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled. For Gigabit Ethernet, when the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

Redundant interfaces (deprecated)

Redundant interfaces were supported for the ASA 5500-X platforms only. We don't recommend configuring them for other platforms.

Default settings for interfaces

This section lists default settings for interfaces.

Default state of interfaces

The default state of an interface depends on the type.

- Physical interfaces—Disabled. The exception is the Management interface that is enabled for initial setup. Physical interfaces includes switch ports.
- VLAN subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.
- EtherChannel port-channel interfaces (ISA 3000)—Enabled. However, for traffic to pass through the EtherChannel, the channel group physical interfaces must also be enabled.
- EtherChannel port-channel interfaces (Firepower and Secure Firewall models)—Disabled.

Note For the Firepower 4100/9300, you can administratively enable and disable interfaces in both the chassis and in the Cloud-Delivered Firewall Management Center. For an interface to be operational, the interface must be enabled in both operating systems. Because the interface state is controlled independently, you may have a mismatch between the chassis and Cloud-Delivered Firewall Management Center.

Default speed and duplex

By default, the speed and duplex for copper (RJ-45) interfaces are set to auto-negotiate.

By default, the speed and duplex for fiber (SFP) interfaces are set to the maximum speed, with auto-negotiation enabled. If a peer switch connecting to the port over a 50G cable does not support auto-negotiation, ensure to disable auto-negotiation on the switch and the Threat Defense interface as well. For example,

N9K-C93400LD-H1 does not support auto-negotiation on a 50G cable. Hence, you must disable the default auto-negotiation on the platform and the switch for the port to be connected.

For the Secure Firewall 3100/4200 the speed is set to detect the installed SFP speed.

Create security zone and interface group objects

Create empty interface objects that can be populated with interfaces later, or create security zones while configuring interfaces to provide logical grouping and management of device interfaces.

Add security zones and interface groups to which you can assign device interfaces.

Tip You can create empty interface objects and add interfaces to them later. To add an interface, the interface must have a name. You can also create security zones (but not interface groups) while configuring interfaces.

Before you begin

Understand the usage requirements and restrictions for each type of interface object. See Security zones and interface groups, on page 999.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Interface**.

You can alternatively choose **Objects** > **Other FTD Objects** to create objects.

Step 2 Click **Add > Security Zone** or **Add > Interface Group**.

Step 3 Enter a **Name**.

Do not use the same name as a network or port object. These object names are deployed to the device, and duplicate names result in a failed deployment.

Step 4 Choose an **Interface Type**.

Step 5 (Optional) From the **Device** > **Interfaces** drop-down list, choose a device that contains interfaces you want to add.

You do not need to assign interfaces on this screen; you can instead assign interfaces to the zone or group when you configure the interface.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Enable the physical interface and configure Ethernet settings

Before you begin

If you changed the physical interfaces on the device after you added it to the Cloud-Delivered Firewall Management Center , you need to refresh the interface listing by clicking **Sync Interfaces from device** on the top left of **Interfaces** . For the 3100 /4200 , which supports hot swapping, see Network module management for Secure Firewall 3100/4200, on page 1015 before you change interfaces on a device.

This section describes how to:

- Enable the physical interface. By default, physical interfaces are disabled (with the exception of the Management interface).
- Set a specific speed and duplex. By default, speed and duplex are set to Auto.

This procedure only covers a small subset of interface settings. Refrain from setting other parameters at this point. For example, you cannot name an interface that you want to use as part of an EtherChannel interface. The exact interface options vary depending on your model and interface type.

Note For the Firepower 4100/9300 , you configure basic interface settings in FXOS. See Configure a physical interface for more information.

Note For switch ports, see Configure switch ports, on page 1040 .

You can also click the **Virtual Tunnels** tab to view details of dynamic and static VTIs of route-based VPNs on the device. For more information, see View virtual tunnel interface details, on page 1769 .

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Enable the interface by checking the **Enabled** check box.

a) (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 4 (Optional) Set the duplex and speed by clicking **Hardware Configuration** > **Speed** .

- **Duplex** —Choose **Full** or **Half** . SFP interfaces only support **Full** duplex. However, the SFP interface (Ethernet 1/5) of Secure Firewall 220 supports **Full** and **Half** duplex.
- **Speed** —Choose a speed (varies depending on the model). For SFPs, choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always full, autonegotiation

is enabled, and FEC is set to auto. For dual-speed transceivers, the lower speed is used. This option is useful if you later change the network module to a different model, and want the speed to update automatically. Some switches and transceivers don't support autonegotiation, especially for higher speed interfaces. In this case, set the interface on the Firewall Threat Defense to a manual speed and also disable **Auto-negotiation** so the link can come up.

Note
You cannot modify the speed of a high availability or a cluster control link interface.

- **Auto-negotiation** —Set the interface to negotiate the link status and flow control.

  Except for the 1100 and 2100 , autonegotiation is set separately from the speed. Some switches don't support autonegotiation, especially for higher speed interfaces. In this case, set the interface on the Firewall Threat Defense to a manual speed and also disable **Auto-negotiation** so the link can come up.

- **Forward Error Correction Mode** —For 25Gbps and higher interfaces, enable Forward Error Correction (FEC).

  For an EtherChannel member interface, configure FEC before adding it to the EtherChannel. If you remove the interface from EtherChannel and then reboot, reconfigure the FEC for the interface.

  Some switches don't support auto mode for FEC, especially for larger interfaces. Be sure to manually configure the setting, depending on the switch support. We don't recommend disabling FEC because the EtherChannel may not function correctly.

  The setting chosen when you use **auto** depends on the transceiver type and whether the interface is fixed (built-in) or on a network module.

Table 66: Default FEC for Auto Setting

Transceiver Type	Fixed Port Default FEC (Ethernet 1/9 through 1/16)	Network Module Default FEC
25G-SR	Clause 74 FC-FEC	Clause 108 RS-FEC
25G-LR	Clause 74 FC-FEC	Clause 108 RS-FEC
10/25G-CSR	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-AOC *x* M	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-CU2.5/3M	Auto-Negotiate	Auto-Negotiate
25G-CU4/5M	Auto-Negotiate	Auto-Negotiate
25/50/100G	Clause 91 RS-FEC	Clause 91 RS-FEC

Note
From Version 7.4.2, Firewall Threat Defense supports the SFP 10G_25G_CSR_S module. To avoid link-down issues after a reboot, Forward Error Correction (FEC) must be enabled on these interfaces.

Step 5 (Optional) Enable Link Layer Discovery Protocol (LLDP) by clicking **Hardware Configuration** > **Network Connectivity** .

- **Enable LLDP Receive** —Enables the firewall to receive LLDP packets from its peers.

- **Enable LLDP Transmit** —Enables the firewall to send LLDP packets to its peers.

Step 6 (Optional) Enable pause (XOFF) frames for flow control by clicking **Hardware Configuration** > **Network Connectivity** , and checking **Flow Control Send** .

Flow control enables connected Ethernet ports to control traffic rates during congestion by allowing congested nodes to pause link operation at the other end. If the threat defense port experiences congestion (exhaustion of queuing resources on the internal switch) and cannot receive any more traffic, it notifies the other port by sending a pause frame to stop sending until the condition clears. Upon receipt of a pause frame, the sending device stops sending any data packets, which prevents any loss of data packets during the congestion period.

Note
The Firewall Threat Defense supports transmitting pause frames so that the remote peer can rate-control the traffic.

However, receiving of pause frames is not supported.

The internal switch has a global pool of 8000 buffers of 250 bytes each, and the switch allocates buffers dynamically to each port. A pause frame is sent out every interface with flowcontrol enabled when the buffer usage exceeds the global high-water mark (2 MB (8000 buffers)); and a pause frame is sent out of a particular interface when its buffer exceeds the port high-water mark (.3125 MB (1250 buffers)). After a pause is sent, an XON frame can be sent when the buffer usage is reduced below the low-water mark (1.25 MB globally (5000 buffers); .25 MB per port (1000 buffers)). The link partner can resume traffic after receiving an XON frame.

Only flow control frames defined in 802.3x are supported. Priority-based flow control is not supported.

Step 7 In the **Mode** drop-down list, choose one of the following:

- **None** —Choose this setting for regular firewall interfaces and inline sets. The mode will automatically be changed to Routed, Switched, or Inline based on further configuration.
- **Passive** —Choose this setting for passive IPS-only interfaces.
- **Erspan** —Choose this setting for ERSPAN passive IPS-only interfaces.

Step 8 In the **Priority** field, enter a number ranging from 0–65535.

This value is used in the policy based routing configuration. The priority is used to determine how you want to distribute the traffic across multiple egress interfaces.

Step 9 Click **OK** .

Click **Save**.

Go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

Continue configuring interfaces.

- Regular Firewall Interfaces, on page 1039
- Inline Sets and Passive Interfaces, on page 1119

Configure EtherChannel interfaces

This section tells how to configure EtherChannel interfaces.

About EtherChannels

This section describes EtherChannels.

EtherChannels

An 802.3ad EtherChannel is a logical interface, called a port-channel interface) that consists of a bundle of individual Ethernet links, known as a channel group. This approach increases the bandwidth for a single network. A port channel interface is used in the same way as a physical interface when you configure interface-related features.

You can configure up to 48 EtherChannels, depending on how many interfaces your model supports.

Channel group interfaces

Channel group interfaces enable EtherChannel aggregation with specific capacity and configuration requirements. Understanding these interface specifications ensures proper EtherChannel deployment and performance.

Each channel group can have up to 8 active interfaces, except for the ISA 3000, which support 16 active interfaces. On switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group. Only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.

All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.

The EtherChannel aggregates traffic across all the available active interfaces in the channel. A proprietary hash algorithm selects the interface, using criteria such as source or destination MAC addresses, IP addresses, TCP and UDP port numbers, and VLAN numbers.

EtherChannel connection to other devices

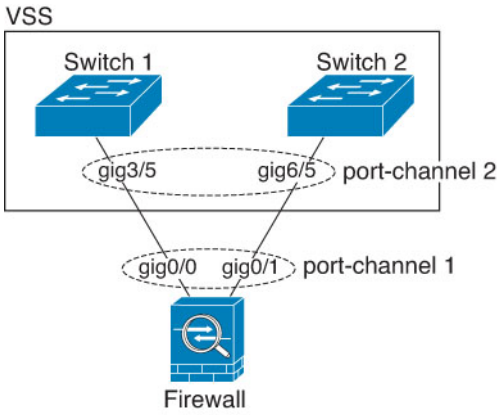
An EtherChannel connection to other devices requires the connected device to support 802.3ad EtherChannels for compatibility.

EtherChannel connection requirements and configurations

The device to which you connect the Firewall Threat Defense EtherChannel must also support 802.3ad EtherChannels. For example, you can connect to the Catalyst 6500 switch or the Cisco Nexus 7000.

When the switch is part of a Virtual Switching System (VSS) or Virtual Port Channel (vPC), you can connect Firewall Threat Defense interfaces within the same EtherChannel to separate switches in the VSS/vPC. The switch interfaces are members of the same EtherChannel port-channel interface because the separate switches act like a single switch.

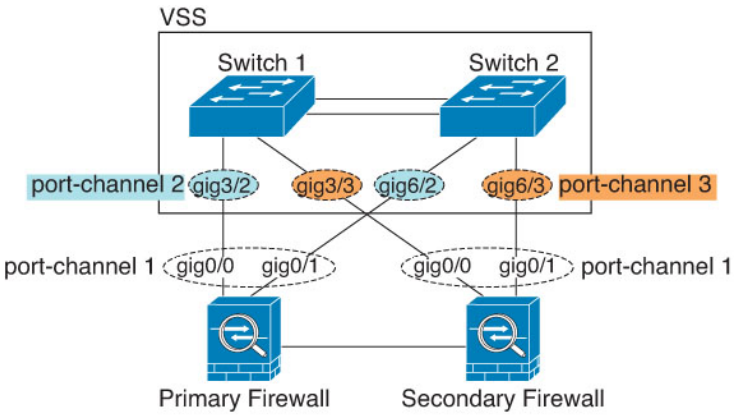
Figure 202: Connecting to a VSS/vPC



Note If the Firewall Threat Defense device is in transparent firewall mode, and you place the Firewall Threat Defense device between two sets of VSS/vPC switches, then be sure to disable Unidirectional Link Detection (UDLD) on any switch ports connected to the Firewall Threat Defense device with an EtherChannel. If you enable UDLD, then a switch port may receive UDLD packets sourced from both switches in the other VSS/vPC pair. The receiving switch will place the receiving interface in a down state with the reason "UDLD Neighbor mismatch".

If you use the Firewall Threat Defense device in an Active/Standby failover deployment, you need to create separate EtherChannels on the switches in the VSS/vPC, one for each Firewall Threat Defense device. On each Firewall Threat Defense device, a single EtherChannel connects to both switches. Even if you could group all switch interfaces into a single EtherChannel connecting to both Firewall Threat Defense devices,the EtherChannel would not be established because of the separate Firewall Threat Defense system IDs. Also, using a single EtherChannel would not be desirable because you do not want traffic sent to the standby Firewall Threat Defense device.

Figure 203: Active/Standby failover and VSS/vPC



Link aggregation control protocol

The Link Aggregation Control Protocol (LACP) is a network protocol that aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

LACP interface configuration modes

You can configure each physical interface in an EtherChannel to be:

- Active—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- Passive—Receives LACP updates. A passive EtherChannel can only establish connectivity with an active EtherChannel. Not supported on hardware models.
- On—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. "On" mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

Load balancing

Load balancing is a distribution mechanism that distributes packets to interfaces in the EtherChannel and ensures transparent failover when active interfaces go down by rebalancing traffic between remaining links.

Load balancing operation details

The Firewall Threat Defense device distributes packets to the interfaces in the EtherChannel by hashing the source and destination IP address of the packet (this criteria is configurable). The resulting hash is divided by the number of active links in a modulo operation where the resulting remainder determines which interface owns the flow. All packets with a *hash_value* **mod** *active_links* result of 0 go to the first interface in the EtherChannel, packets with a result of 1 go to the second interface, packets with a result of 2 go to the third interface, and so on. For example, if you have 15 active links, then the modulo operation provides values from 0 to 14. For 6 active links, the values are 0 to 5, and so on.

If an active interface goes down and is not replaced by a standby interface, then traffic is rebalanced between the remaining links. The failure is masked from both Spanning Tree at Layer 2 and the routing table at Layer 3, so the switchover is transparent to other network devices.

EtherChannel MAC address

An EtherChannel MAC address is a network identifier that is shared by all interfaces in the EtherChannel group. This makes the EtherChannel transparent to network applications and users by presenting one logical connection.

All interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links.

Platform-specific MAC address behavior

MAC address assignment for EtherChannel interface varies by hardware platform.

Firepower and Secure Firewall Hardware

The port-channel interface uses the MAC address of the internal interface Internal-Data 0/1. You can manually configure a MAC address for the port-channel interface if needed. All EtherChannel interfaces on a chassis use the same MAC address, so be aware that if you use SNMP polling, for example, multiple interfaces will have the same MAC address.

Note Member interfaces only use the Internal-Data 0/1 MAC address after a reboot. Prior to rebooting, the member interface uses its own MAC address. If you add a new member interface after a reboot, you will have to perform another reboot to update its MAC address.

Guidelines for EtherChannels and redundant interfaces

Provides essential configuration rules and limitations for implementing EtherChannels and redundant interfaces to maintain network stability and high availability.

Bridge Group

In routed mode, Cloud-Delivered Firewall Management Center -defined EtherChannels are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.

High availability

- When you use an EtherChannel interface as a High availability link, it must be pre-configured on both units in the High availability pair; you cannot configure it on the primary unit and expect it to replicate to the secondary unit because the High availability link itself is required for replication .
- If you use an EtherChannel interface for the state link, no special configuration is required; the configuration can replicate from the primary unit as normal. For the Firepower 4100/9300 chassis , all interfaces, including EtherChannels, need to be pre-configured on both units.
- You can monitor EtherChannel interfaces for High availability . When an active member interface fails over to a standby interface, this activity does not cause the EtherChannel interface to appear to be failed when being monitored for device-level High availability . Only when all physical interfaces fail does the EtherChannel interface appear to be failed (for an EtherChannel interface, the number of member interfaces allowed to fail is configurable) .
- If you use an EtherChannel interface for a High availability or state link, then to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a High availability link. To alter the configuration, you need to temporarily disable High availability , which prevents High availability from occurring for the duration.

Model Support

- You cannot add EtherChannels in the Cloud-Delivered Firewall Management Center for the Firepower 4100/9300 or the Firewall Threat Defense Virtual . The Firepower 4100/9300 supports EtherChannels, but you must perform all hardware configuration of EtherChannels in FXOS on the chassis.

- You cannot use Firepower 1010 or Secure Firewall 1210/1220 switch ports or VLAN interfaces in EtherChannels.

Clustering

- To configure a spanned EtherChannel or an individual cluster interface, see the clustering chapter.

General Redundant Interface Guidelines

- You can configure up to 8 redundant interface pairs.
- All the Firewall Threat Defense configuration refers to the logical redundant interface instead of the member physical interfaces.
- You cannot use a redundant interface as part of an EtherChannel, nor can you use an EtherChannel as part of a redundant interface. You cannot use the same physical interfaces in a redundant interface and an EtherChannel interface. You can, however, configure both types on the Firewall Threat Defense device if they do not use the same physical interfaces.
- If you shut down the active interface, then the standby interface becomes active.
- Redundant interfaces do not support the Management *slot* / *port* interfaces as members. You can, however, set a redundant interface comprised of non- Management interfaces as management-only.

General EtherChannel Guidelines

- You can configure up to 48 EtherChannels, depending on how many interfaces are available on your model.
- Each channel group can have up to 8 active interfaces, except for ASA models and the ISA 3000, which supports 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- When you add the first member interface, it sets the required hardware properties of all member interfaces.
    - The media type of member interfaces can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix RJ-45 and SFP interfaces.
    - All interfaces must be set to the same speed and duplex.
    - The first interface sets the speed *capacity* , which cannot be changed later.
- The device to which you connect the Firewall Threat Defense EtherChannel must also support 802.3ad EtherChannels.
- The Firewall Threat Defense device does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **vlan dot1Q tag native** command, then the Firewall Threat Defense device will drop the tagged LACPDUs.
- The LACP rate depends on the model. When you set the rate (normal or fast), the device requests that rate from the connecting switch.

- In Cisco IOS software versions earlier than 15.1(1)S2, Firewall Threat Defense does not support connecting an EtherChannel to a switch stack. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time.
- Match the speed, duplex, and auto-negotiation settings on every EtherChannel member interface and its corresponding peer port. For a 10/100-Mbps member link, use auto-negotiation at both ends, or use the same fixed speed and duplex settings at both ends. If one end is fixed at full duplex while the other end uses auto-negotiation, the auto-negotiating side defaults to half duplex, resulting in a duplex mismatch.
- All the Firewall Threat Defense configuration refers to the logical EtherChannel interface instead of the member physical interfaces.

Configure an EtherChannel

Create an EtherChannel port-channel interface to aggregate multiple physical interfaces into a single logical interface for increased bandwidth and redundancy.

This section describes how to create an EtherChannel port-channel interface, assign interfaces to the EtherChannel, and customize the EtherChannel.

Note For the Firepower 4100/9300, you configure EtherChannels in FXOS. See Add an EtherChannel (port channel) for more information.

Before you begin

- When you add the first member interface, it sets the required hardware properties of all member interfaces. See Guidelines for EtherChannels and redundant interfaces, on page 1009 for more details about member interface requirements.
- You cannot add a physical interface to the channel group if you configured a name for it. You must first remove the name.

Note If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Follow these steps to configure an EtherChannel:

Procedure

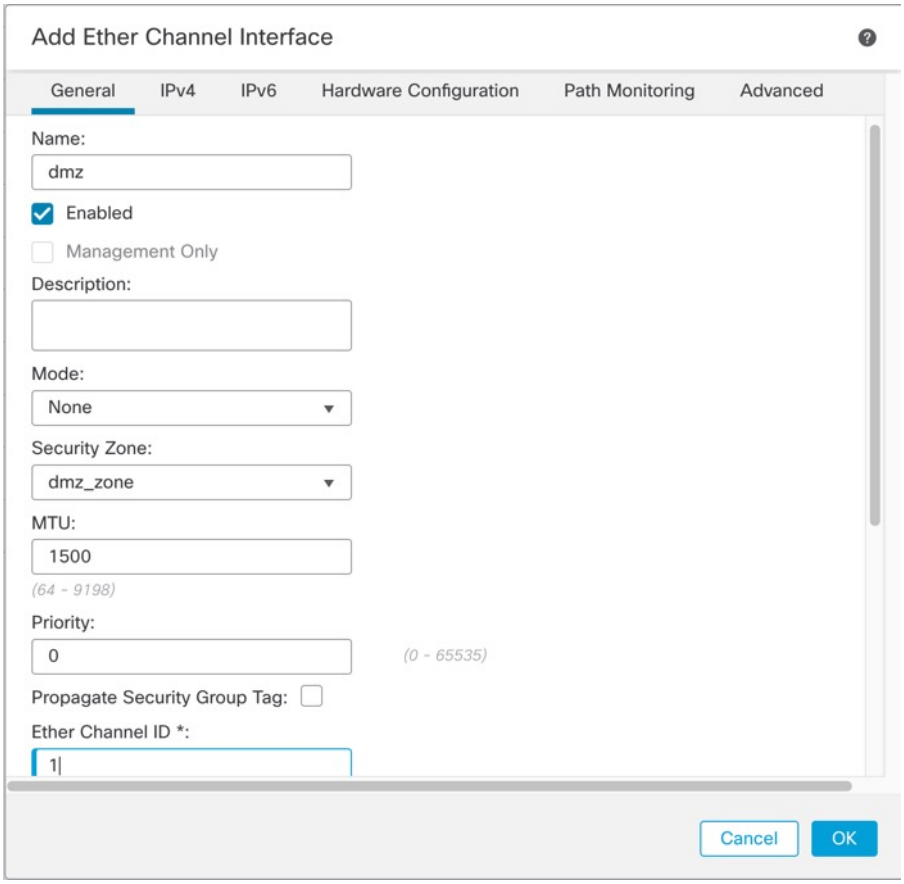
Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the member interfaces according to Enable the physical interface and configure Ethernet settings, on page 1003.

Step 3 Click **Add Interfaces** > **Ether Channel Interface**.

Step 4 On the **General** tab, set the **Ether Channel ID** to a number between 1 and 48 (1 and 8 for the Firepower 1010 and Secure Firewall 1210, 1 and 10 for the Secure Firewall 1220).

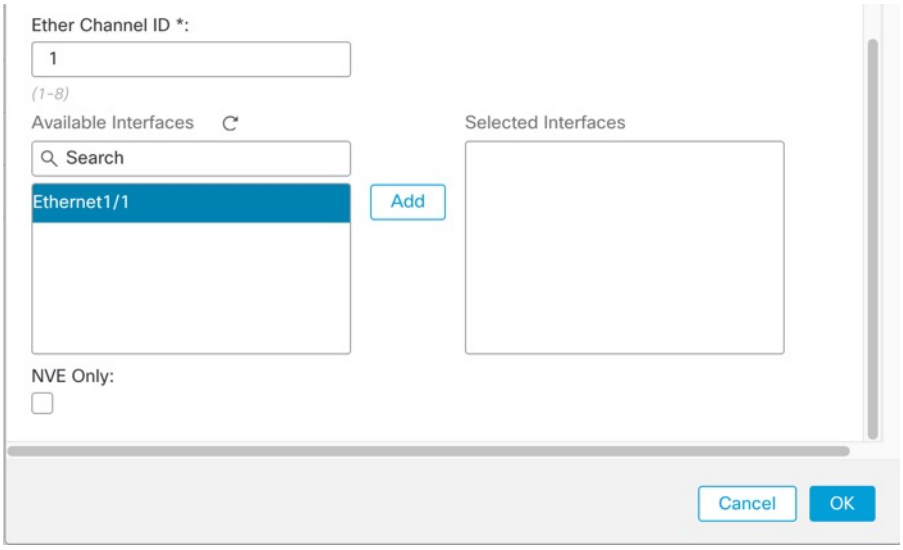
Figure 204: Add EtherChannel interface



Step 5 In the **Available Interfaces** area, click an interface and then click **Add** to move it to the **Selected Interfaces** area. Repeat for all interfaces that you want to make members.

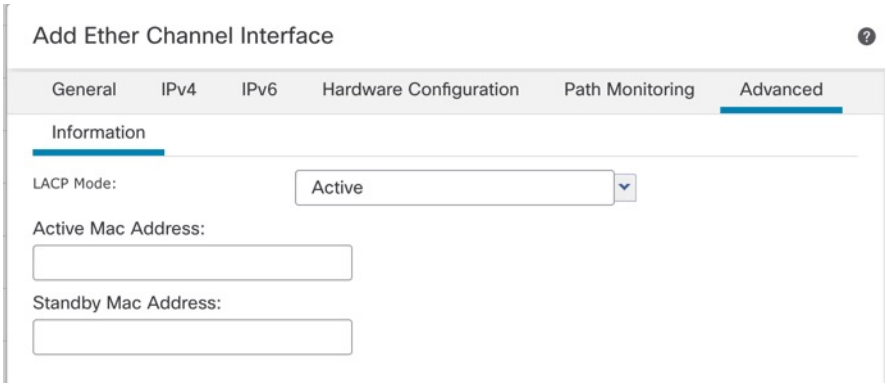
Make sure all interfaces are the same type and speed capability.

Figure 205: Available interfaces



Step 6 (Optional) Click the **Advanced** tab to customize the EtherChannel. Set the following parameters on the **Information** sub-tab:

Figure 206: Advanced



- (ISA 3000 only) **Load Balancing**—Select the criteria used to load balance the packets across the group channel interfaces. By default, the Firewall Threat Defense device balances the packet load on interfaces according to the source and destination IP address of the packet. If you want to change the properties on which the packet is categorized, choose a different set of criteria. For example, if your traffic is biased heavily towards the same source and destination IP addresses, then the traffic assignment to interfaces in the EtherChannel will be unbalanced. Changing to a different algorithm can result in more evenly distributed traffic. For more information about load balancing, see Load balancing, on page 1008.
- **LACP Mode**—Choose Active, Passive, or On. We recommend using Active mode (the default). Passive mode is only available for the ISA 3000 only.
- (Secure Firewall 3100/4200 only) **LACP Rate**—Choose Default, Normal, or Fast. The defualt is Normal (also known as slow). Sets the LACP data unit receive rate for a physical interface in the channel group. We recommend that you set the same rate on both sides.

- (ISA 3000 only) **Active Physical Interface: Range**—From the left drop-down list, choose the minimum number of active interfaces required for the EtherChannel to be active, between 1 and 16. The default is 1. From the right drop-down list, choose the maximum number of active interfaces allowed in the EtherChannel, between 1 and 16. The default is 16. If your switch does not support 16 active interfaces, be sure to set this command to 8 or fewer.
- **Active MAC Address**—Set a manual MAC address if desired. The mac_address is in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE is entered as 000C.F142.4CDE.

Step 7 Click the **Hardware Configuration** tab and set the Duplex and Speed for all member interfaces.

Step 8 Click **OK**.

Step 9 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

For regular firewall interfaces:

- Add a VLAN subinterface. See Add a subinterface, on page 1060.
- Configure the routed or transparent mode interface parameters: Configure routed mode interfaces, on page 1080 or Bridge group interfaces, on page 1086. For IPS-only interfaces, see Inline Sets and Passive Interfaces, on page 1119.

Sync Interface Changes with the Cloud-Delivered Firewall Management Center

Synchronize interface changes to ensure the Cloud-Delivered Firewall Management Center and device configurations remain aligned when physical interfaces are added or removed.

Addition or deletion of physical interfaces on the device can cause the Cloud-Delivered Firewall Management Center and the device to get out of sync. The Cloud-Delivered Firewall Management Center can detect interface changes by one of these methods:

- Event sent from the device
- Synchronize when you deploy from the Cloud-Delivered Firewall Management Center

  If the Cloud-Delivered Firewall Management Center detects interface changes when it attempts to deploy, the deployment will fail. You must first accept the interface changes.
- Manual synchronization

Adding a new interface, or deleting an unused interface has minimal impact on the Firewall Threat Defense configuration. However, deleting an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the Firewall Threat Defense configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any

configuration associated with that interface. Policies that refer to security zones are not affected. You can also edit the membership of an allocated EtherChannel without affecting the logical device or requiring a sync on the Cloud-Delivered Firewall Management Center.

When the Cloud-Delivered Firewall Management Center detects changes, the **Interface** page shows status (removed, changed, or added) to the left of each interface.

This procedure describes how to manually sync interface changes if required. If interface changes are temporary, you should not save the changes in the Cloud-Delivered Firewall Management Center; you should wait until the device is stable, and then re-sync.

Before you begin

Follow these steps to sync interface changes with the Cloud-Delivered Firewall Management Center:

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 If required, click **Sync Interfaces** on the top left of **Interfaces**.

Step 3 After the changes are detected, see these steps.

a) You will see a red banner on **Interfaces** indicating that the interface configuration has changed. Click the **Click to know more** link to view the interface changes.

b) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

c) Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices.

Step 4

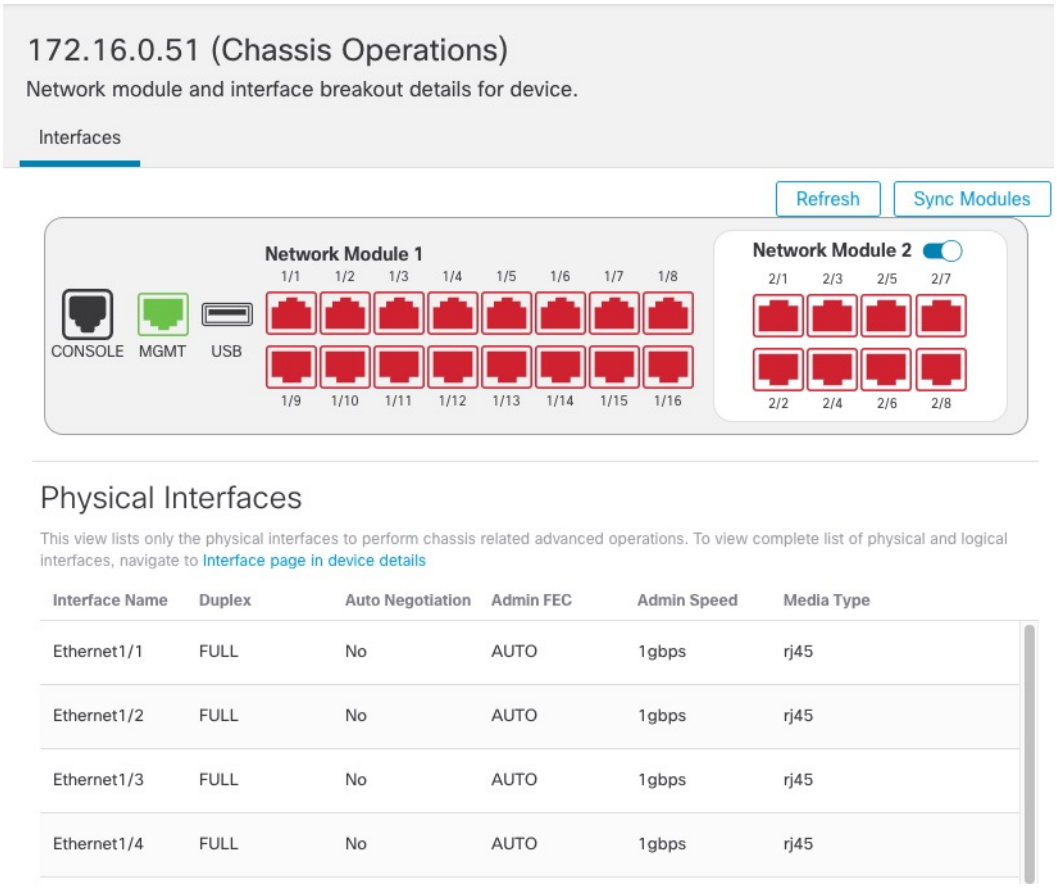
Network module management for Secure Firewall 3100/4200

If you install a network module before powering on the device, the network module becomes enabled and ready for use automatically.

Network module management operations

To view physical interface details for the device, and to manage the network module, open the **Chassis Operations** page. From **Devices** > **Device Management**, click **Manage** in the **Chassis** column. For clustering or High Availability, this option is only available for the control node/active unit. The **Chassis Operations** page opens for the device.

Figure 207: Chassis operations



Click **Refresh** to refresh interface status. Click **Sync Modules** if you made a hardware change on the device that you need to detect.

If you need to make changes to your network module installation after initial bootup, then see the following procedures.

Configure breakout ports

This task configures breakout ports to create smaller ports from 40GB or higher interfaces, allowing you to split a single high-capacity interface into multiple physical Ethernet ports that can be used like any other physical Ethernet port, including being added to EtherChannels.

You can configure smaller breakout ports for each 40GB or higher interface. This procedure tells you how to break out and rejoin the ports. The breakout ports can be used just like any other physical Ethernet port, including being added to EtherChannels. For example, you can break out four 10GB ports from a 40GB interface. The exact size and number of ports depends on your model.

Changes are immediate; you do not need to deploy to the device. After you break or rejoin, you cannot roll back to the previous interface state.

Before you begin

- You must use a supported breakout cable. See the hardware installation guide for more information.

- The interface cannot be in use for the following before breaking or rejoining:
  - Failover link
  - Cluster control link
  - Have a subinterface
  - EtherChannel member
  - BVI member
  - Manager access interface
- Breaking or rejoining an interface that is used directly in your security policy can impact the configuration; however, the action is not blocked.

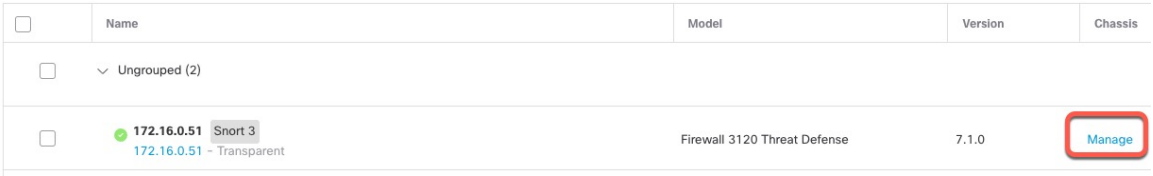
Follow these steps to configure breakout ports:

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column.

For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

Figure 208: Manage chassis



The **Chassis Operations** page opens for the device (in multi-instance mode, this page is called **Chassis Manager**). This page shows physical interface details for the device.

Step 2 Break out ports from a 40GB or higher interface.

a) Click **Break** ( ) to the right of the interface.

Click **Yes** on the confirmation dialog box. If the interface is in use, you will see an error message. You must resolve any use cases before you can retry the breakout.

For example, to break out the Ethernet2/1 40GB interface, the resulting child interfaces will be identified as Ethernet2/1/1, Ethernet2/1/2, Ethernet2/1/3, and Ethernet2/1/4.

On the interfaces graphic, a port that is broken out has this appearance:

Figure 209: Breakout ports



b) Click the link in the message at the top of the screen to go to the **Interfaces** page to save the interface changes.

Figure 210: Go to interface page



c) At the top of the **Interfaces** page, click **Click to know more**.

The **Interface Changes** dialog box opens.

Figure 211: View interface changes

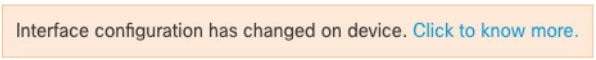
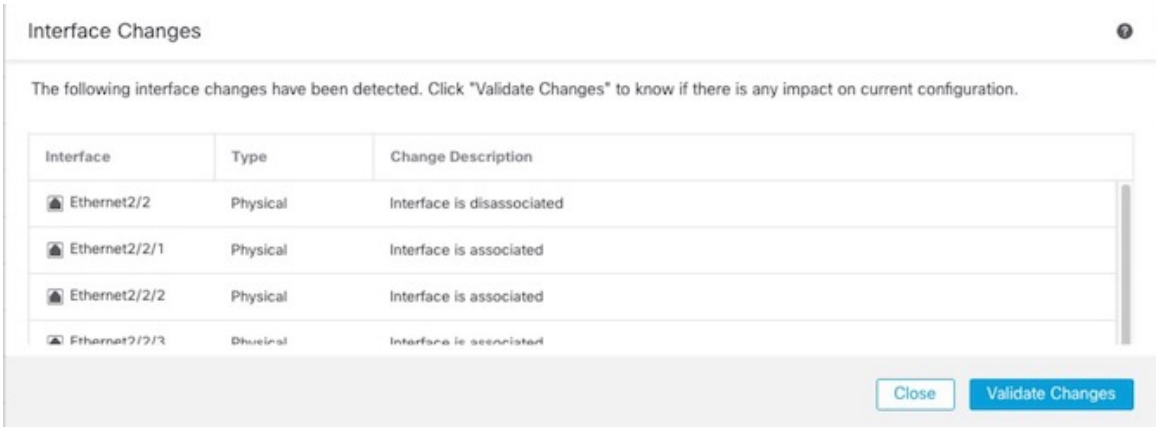


Figure 212: Interface changes



d) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, change your policy and rerun the validation.

Replacing the parent interface that is used in your security policy can impact the configuration. Interfaces can be referenced directly in many places in the configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected.

e) Click **Close** to return to the **Interfaces** page.

f) Click **Save** to save the interface changes to the firewall.

g) If you had to change any configuration, go to **Deploy** > **Deployment** and deploy the policy.

You do not need to deploy just to save the breakout port changes.

Step 3 Rejoin breakout ports.

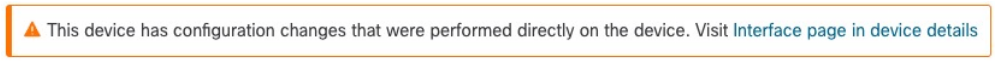
You must rejoin all child ports for the interface.

a) Click **Join** ( ) to the right of the interface.

Click **Yes** on the confirmation dialog box. If any child ports are in use, you will see an error message. You must resolve any use cases before you can retry the rejoin.

b) Click the link in the message at the top of the screen to go to the **Interfaces** page to save the interface changes.

Figure 213: Go to interface page



c) At the top of the **Interfaces** page, click **Click to know more**.

The **Interface Changes** dialog box opens.

Figure 214: View interface changes

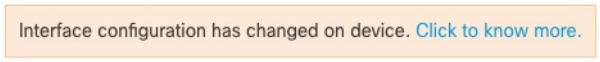
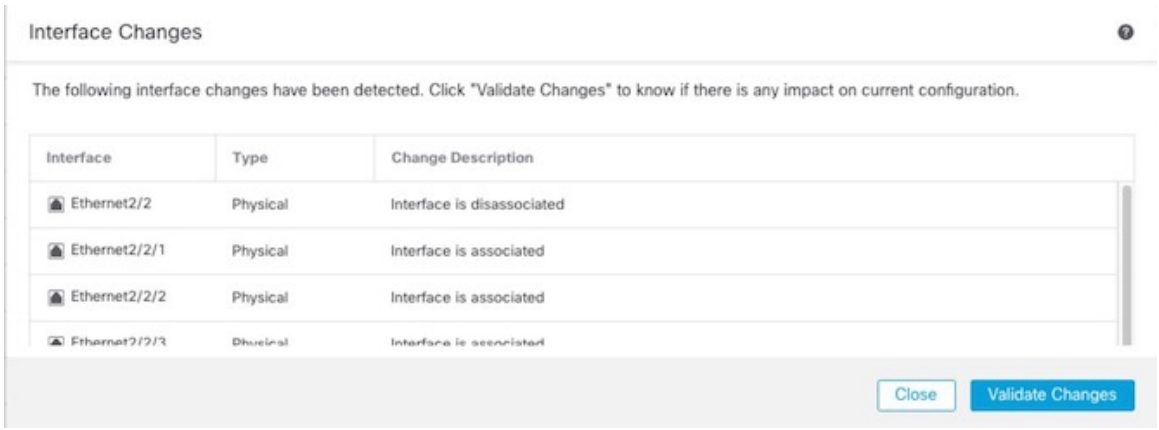


Figure 215: Interface changes



d) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

Replacing the child interfaces that are used in your security policy can impact the configuration. Interfaces can be referenced directly in many places in the configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected.

e) Click **Close** to return to the **Interfaces** page.

f) Click **Save** to save the interface changes to the firewall.

g) If you had to change any configuration, go to **Deploy** > **Deployment** and deploy the policy.

You do not need to deploy just to save the breakout port changes.

Add a network module

This task enables you to add network module functionality to a firewall that is already operational, expanding its interface capabilities.

Adding a new module requires a reboot of the firewall system.

Procedure

Step 1 Install the network module according to the hardware installation guide.

For clustering or High Availability, install the network module on all nodes.

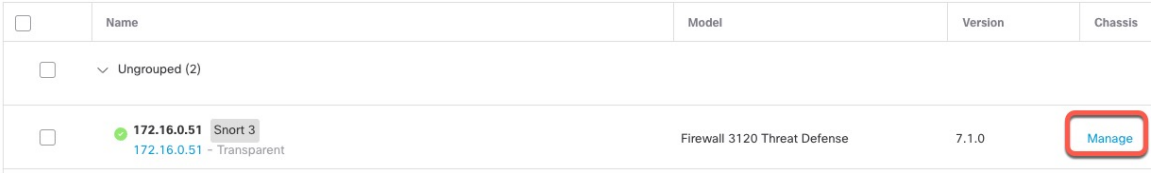
Step 2 Reboot the firewall; see Shut down or restart the device, on page 107.

For clustering or High Availability, reboot the data nodes/standby unit first, and wait for them to come back up. Then you can change the control node or active unit (see Switch the Active Peer in the Firewall Threat Defense High Availability Pair, on page 710), and reboot the former control node/active unit.

Step 3 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column.

For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

Figure 216: Manage chassis

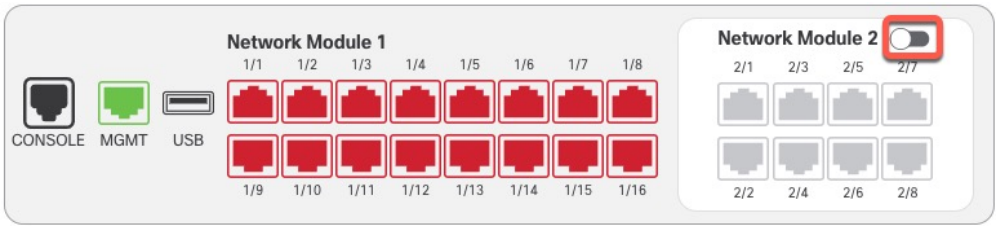


The **Chassis Operations** page opens for the device. This page shows physical interface details for the device.

Step 4 Click **Sync Modules** to update the page with the new network module details.

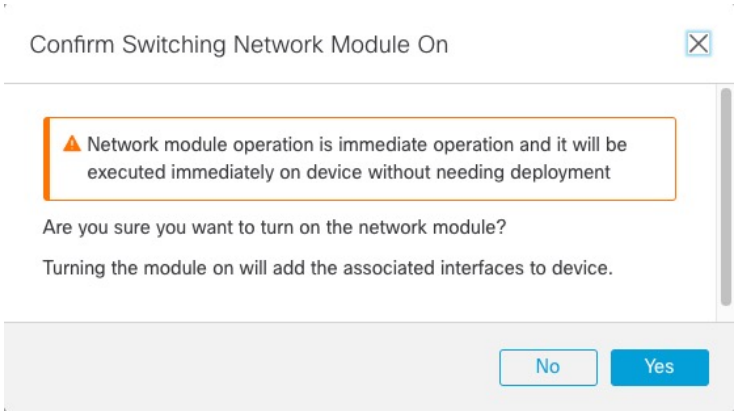
Step 5 On the interfaces graphic, click the slider ( ) to enable the network module.

Figure 217: Enable the network module



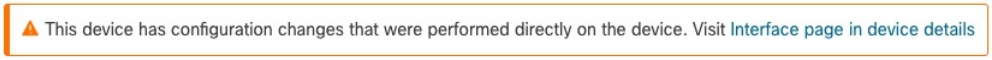
Step 6 You are prompted to confirm that you want to turn the network module on. Click **Yes**.

Figure 218: Confirm enable



Step 7 You see a message at the top of the screen; click the link to go to the **Interfaces** page to save the interface changes.

Figure 219: Go to interface page



Step 8 (Optional) At the top of the **Interfaces** page, you see a message that the interface configuration has changed. You can click **Click to know more** to open the **Interface Changes** dialog box to view the changes.

Figure 220: View interface changes

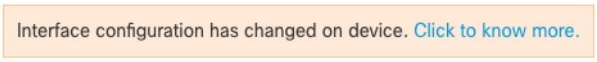
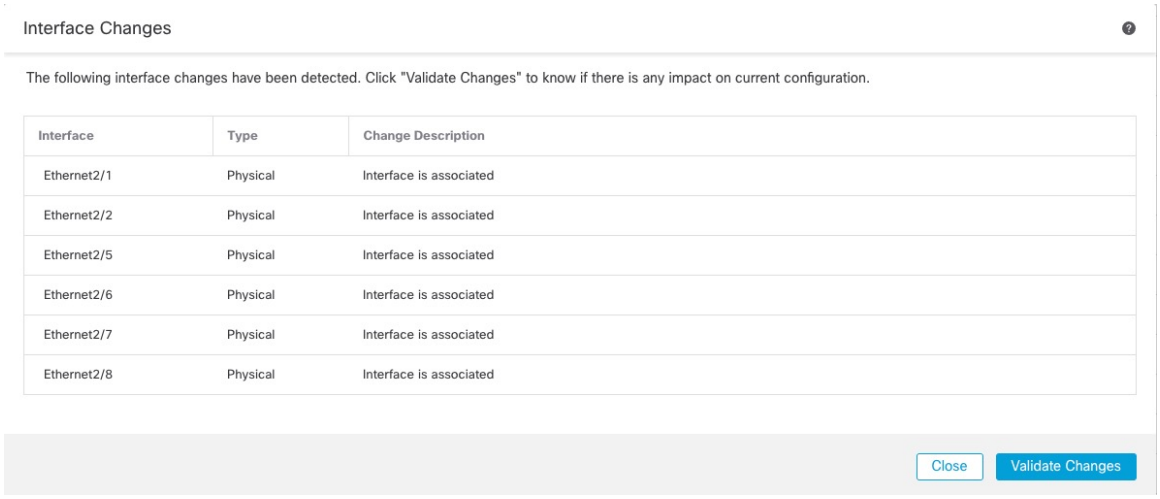


Figure 221: Interface changes



Click **Close** to return to the **Interfaces** page. (Because you are adding a new module, there shouldn't be any configuration impact, so you do not need to click **Validate Changes**.)

Step 9 Click **Save** to save the interface changes to the firewall.

Hot swap the network module

You can hot swap a network module for a new module of the same type without having to reboot.

You must shut down the current module to remove it safely. This procedure describes how to shut down the old module, install a new module, and enable it.

For clustering or High Availability, you can only perform chassis operations on the control node/active unit. You cannot disable a network module if the cluster control link/failover link is on the module.

Procedure

Step 1 For clustering or High Availability, perform the following steps.

- **Clustering**—Ensure the unit you want to perform the hot swap on is a data node; then break the node so it is no longer in the cluster.

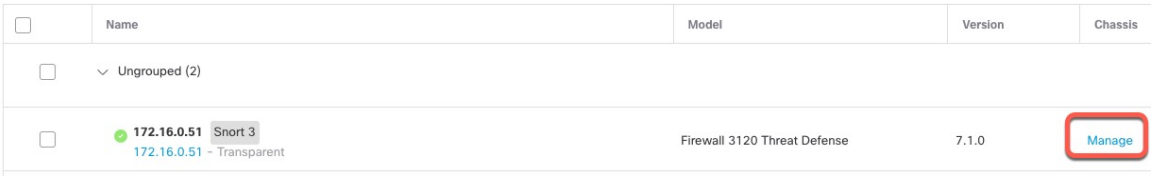
  You will add the node back to the cluster after you perform the hot swap. Alternatively, you can perform all operations on the control node, and the network module changes will sync to all data nodes. However, you will lose use of those interfaces on all nodes during the hot swap.

- **High Availability**—To avoid failing over when you disable the network module:
  - If the failover link is on the network module, you must break High Availability. See Break a high availability pair, on page 714. Disabling the network module with an active failover link is not allowed.
  - Disable interface monitoring for interfaces on the network module. See Configure standby IP addresses and interface monitoring, on page 708.

Step 2 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column.

For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

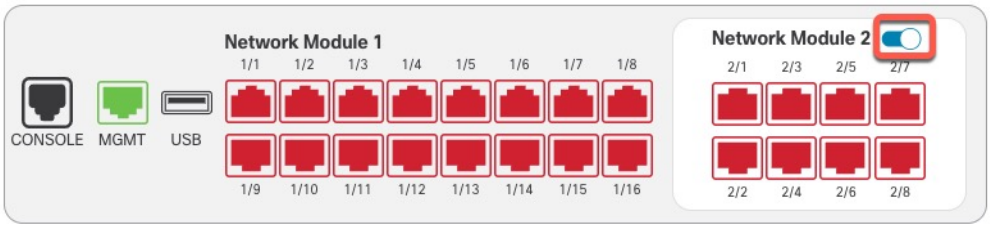
Figure 222: Manage chassis



The **Chassis Operations** page opens for the device. This page shows physical interface details for the device.

Step 3 On the interfaces graphic, click the slider ( ) to disable the network module.

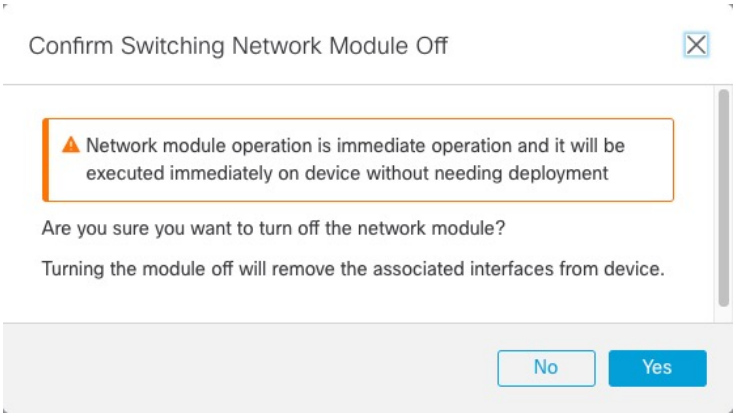
Figure 223: Disable the network module



Do not save any changes on the **Interfaces** page. Because you are replacing the network module, you do not want to disrupt any existing configuration.

Step 4 You are prompted to confirm that you want to turn the network module off. Click **Yes**.

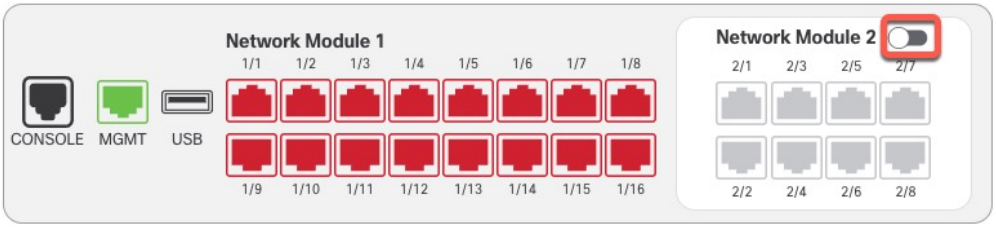
Figure 224: Confirm disable



Step 5 On the device, remove the old network module and replace it with the new network module according to the hardware installation guide.

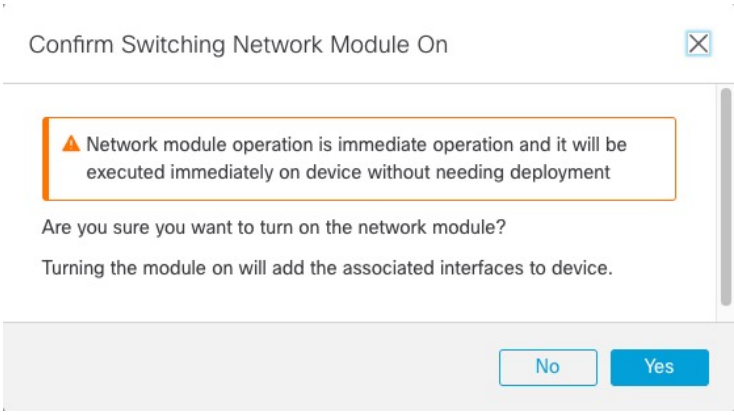
Step 6 In the Cloud-Delivered Firewall Management Center, enable the new module by clicking the slider ().

Figure 225: Enable the network module



Step 7 You are prompted to confirm that you want to turn the network module on. Click **Yes**.

Figure 226: Confirm enable



Step 8 For clustering or High Availability, perform the following steps.

- **Clustering**—Add the node back to the cluster.
- **High Availability**—
  - If you broke High Availability, then reform High Availability. See Add a high availability pair, on page 705.
  - Reenable interface monitoring for interfaces on the network module. See Configure standby IP addresses and interface monitoring, on page 708.

Replace the network module with a different type

Replace an existing network module with a different type to accommodate new interface requirements or hardware upgrades.

If you replace a network module with a different type, then a reboot is required. If the new module has fewer interfaces than the old module, you will have to manually remove any configuration related to interfaces that will no longer be present.

For clustering or High Availability, you can only perform chassis operations on the control node or active unit.

Before you begin

For High Availability, you cannot disable a network module if the failover link is on the module. You will have to break the high availability pair (see Break a high availability pair, on page 714), which means you will have downtime when you reboot the active unit. After the units finish rebooting, you can reform High Availability.

Follow these steps to replace the network module with a different type:

Procedure

Step 1 For clustering or High Availability, perform the following steps.

- **Clustering**—To avoid downtime, you can break each node one at a time so it is no longer in the cluster when you perform the network module replacement.

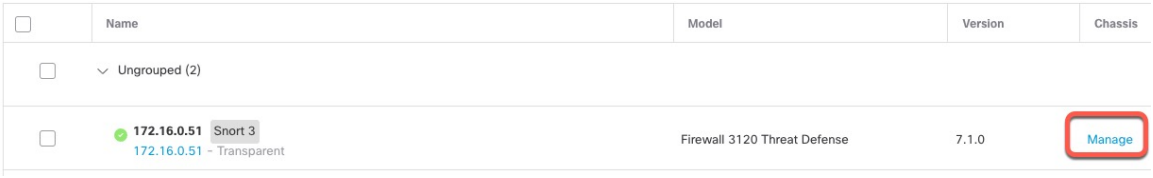
  Add the node back to the cluster after the replacement.

- **High Availability**—To avoid failing over when you replace the network module, disable interface monitoring for interfaces on the network module. See Configure standby IP addresses and interface monitoring, on page 708.

Step 2 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column.

For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

Figure 227: Manage chassis



The **Chassis Operations** page opens for the device. This page shows physical interface details for the device.

Step 3 On the interfaces graphic, click the slider ( ) to disable the network module.

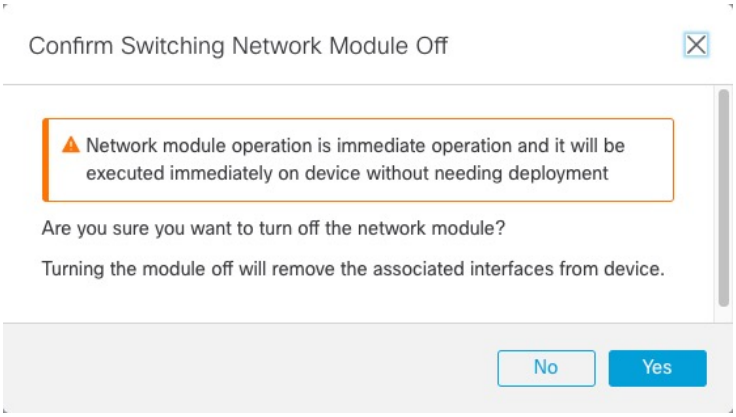
Figure 228: Disable the network module



Do not save any changes on the **Interfaces** page. Because you are replacing the network module, you do not want to disrupt any existing configuration.

a) You are prompted to confirm that you want to turn the network module off. Click **Yes**.

Figure 229: Confirm disable



Step 4 On the device, remove the old network module and replace it with the new network module according to the hardware installation guide.

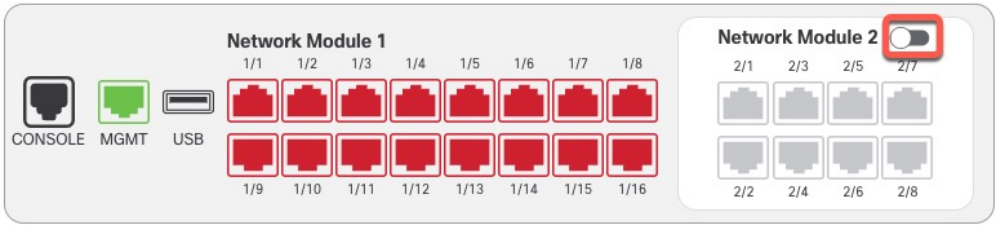
Reboot the firewall. Refer to Shut down or restart the device, on page 107 for guidance.

For clustering or High Availability, reboot the data nodes/standby unit first, and wait for them to come back up. Then you can change the control node or active unit (see Switch the Active Peer in the Firewall Threat Defense High Availability Pair, on page 710), and reboot the former control node/active unit.

Step 5 In the Cloud-Delivered Firewall Management Center, click **Sync Modules** to update the page with the new network module details.

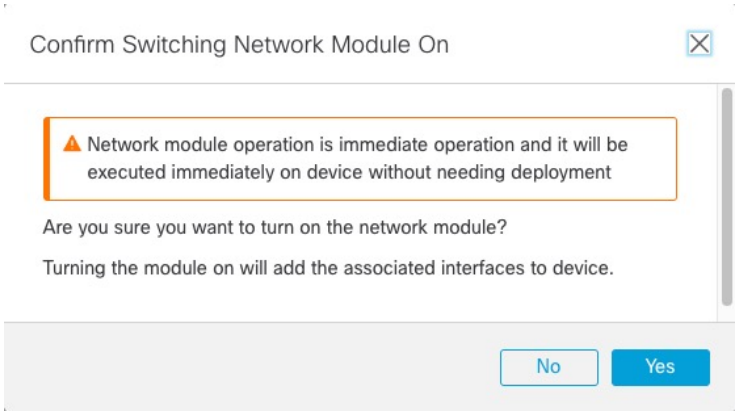
Step 6 Enable the new module by clicking the slider ( ).

Figure 230: Enable the network module



a) Click **Yes** when prompted to confirm that you want to turn the network module on.

Figure 231: Confirm enable



Step 7 Click the link in the message at the top of the screen to go to the **Interfaces** page to save the interface changes.

Figure 232: Go to interface page



Step 8 If the network module has *fewer* interfaces:

a) At the top of the **Interfaces** page, click **Click to know more**.

The **Interface Changes** dialog box opens.

Figure 233: View interface changes

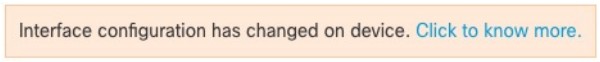
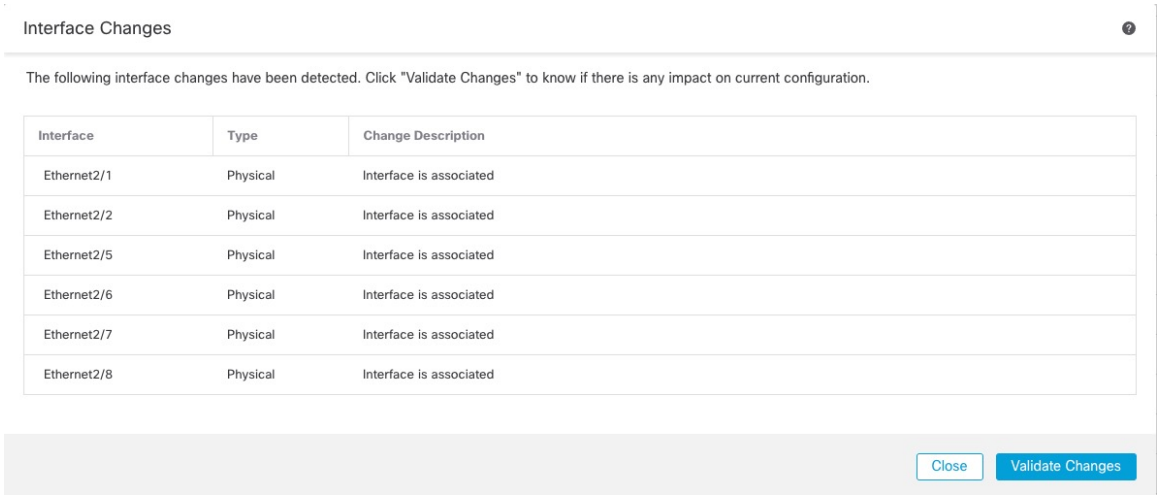


Figure 234: Interface changes



b) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

Deleting an interface that is used in your security policy can impact the configuration. Interfaces can be referenced directly in many places in the configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected.

c) Click **Close** to return to the **Interfaces** page.

Step 9 To change the interface speed, see Enable the physical interface and configure Ethernet settings, on page 1003.

The default speed is set to Detect SFP, which detects the correct speed from the SFP installed. You only need to fix the speed if you manually set the speed to a particular value and you now need a new speed.

Click **Save** to save the interface changes to the firewall.

What to do next

- If you had to change any configuration, go to **Deploy** > **Deployment** and deploy the policy.
- For clustering, add the node back to the cluster.
- For high availability, reenable interface monitoring for interfaces on the network module. See Configure standby IP addresses and interface monitoring, on page 708.

Remove the network module

Remove a network module permanently from the chassis when it is no longer needed or requires replacement.

If you want to permanently remove the network module, follow these steps. Removing a network module requires a reboot.

For clustering or High Availability, you can only perform chassis operations on the control node or active unit.

Before you begin

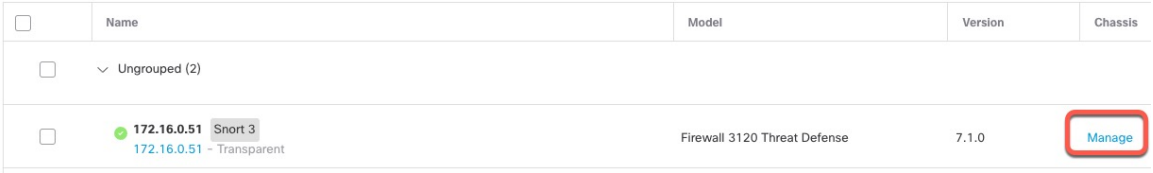
For clustering or High Availability, make sure the cluster or failover link is not on the network module.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column.

For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

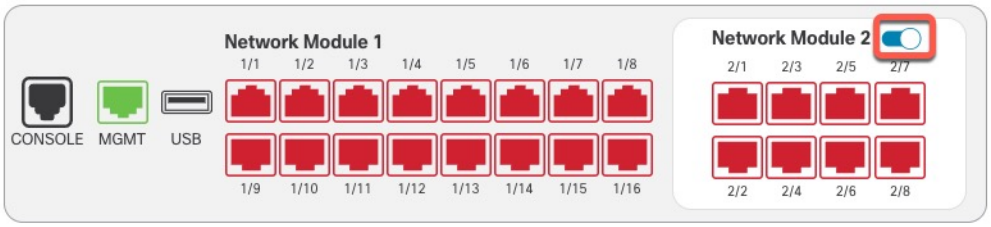
Figure 235: Manage chassis



The **Chassis Operations** page opens for the device. This page shows physical interface details for the device.

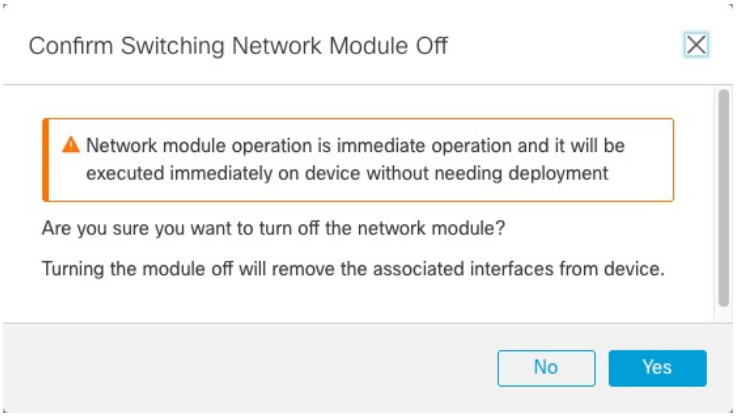
Step 2 On the interfaces graphic, click the slider ( ) to disable the network module.

Figure 236: Disable the network module



Step 3 You are prompted to confirm that you want to turn the network module off. Click **Yes**.

Figure 237: Confirm disable



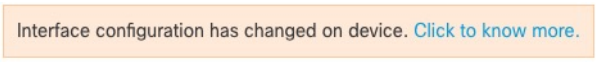
Step 4 You see a message at the top of the screen; click the link to go to the **Interfaces** page to save the interface changes.

Figure 238: Go to interface page



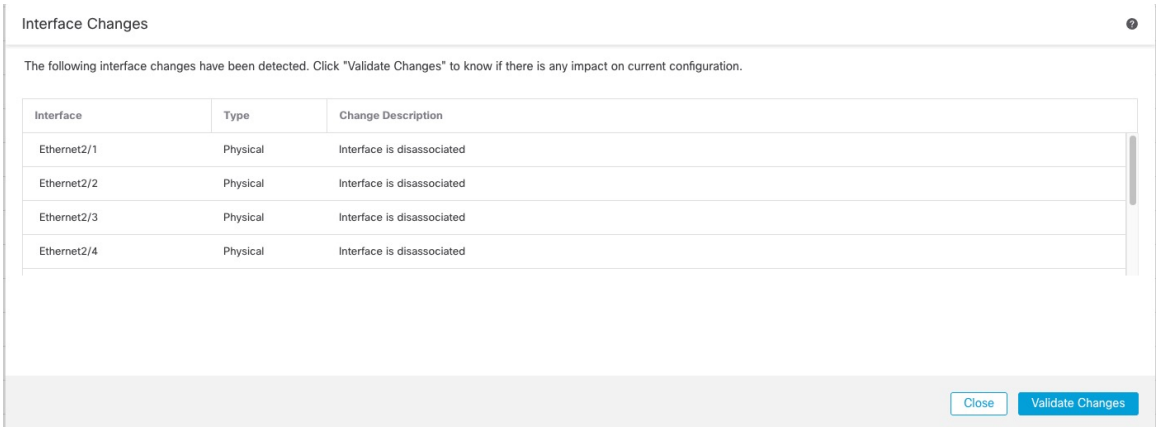
Step 5 At the top of the **Interfaces** page, you see a message that the interface configuration has changed.

Figure 239: View interface changes



a) Click **Click to know more** to open the **Interface Changes** dialog box to view the changes.

Figure 240: Interface changes



b) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

Deleting an interface that is used in your security policy can impact the configuration. Interfaces can be referenced directly in many places in the configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected.

c) Click **Close** to return to the **Interfaces** page.

Step 6 Click **Save** to save the interface changes to the firewall.

Step 7 If you had to change any configuration, go to **Deploy** > **Deployment** and deploy the policy.

Step 8 Reboot the firewall; see Shut down or restart the device, on page 107.

For clustering or High Availability, reboot the data nodes/standby unit first, and wait for them to come back up. Then you can change the control node or active unit (see Switch the Active Peer in the Firewall Threat Defense High Availability Pair, on page 710), and reboot the former control node/active unit.

Merge the management and diagnostic interfaces

This task merges the Management and Diagnostic interfaces on devices that have existing Diagnostic interface configurations that prevent automatic merging.

Starting with version 7.4, Firewall Threat Defense supports a merged Management and Diagnostic interface. "If your configuration uses the Diagnostic interface, the interfaces are not merged automatically. You must perform this procedure to complete the merge. This procedure requires you to acknowledge configuration changes and, in some cases, manually fix the configuration.

The Backup/Restore and Cloud-Delivered Firewall Management Center configuration rollback functions save and restore the merged state, whether it is non-merged or merged. For example, if you merge the interfaces and then restore an old non-merged configuration, the restored configuration will be in a non-merged state.

This table shows the available configuration on the legacy Diagnostic interface, and how the merge is completed.

Before you begin

- To view the current mode of the device, enter the **show management-interface convergence** command at the Firewall Threat Defense CLI. The following output shows that the Management interfaces are merged:

```
> show management-interface convergence
management-interface convergence
>
```

  The following output shows that the Management interfaces are not merged:

```
> show management-interface convergence
no management-interface convergence
>
```

- For High Availability pairs and clusters, perform this task on the active/control unit. The merged configuration will be replicated automatically to the standby/data units.

Follow these steps to merge the Management and Diagnostic interfaces:

Procedure

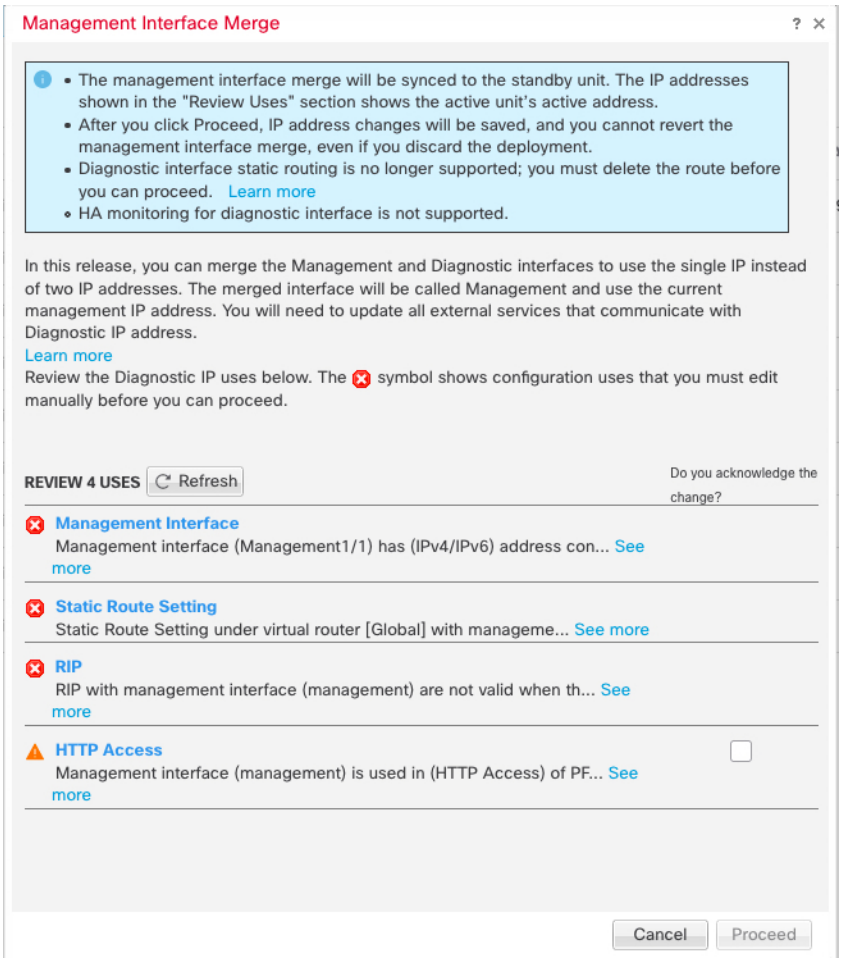
Step 1 Choose **Devices** > **Device Management**, and click **Edit** ( ) for your Firewall Threat Defense. The **Interfaces** page is selected by default. .

Step 2 Edit the Diagnostic interface, and remove the IP address.

You cannot complete the merge until after you have removed the Diagnostic IP address.

Step 3 Click **Management Interface Merge** in the **Management Interface action needed** area.

The **Management Interface Merge** dialog box shows all the occurrences of the Diagnostic interface in the configuration. For any occurrences that require you to manually remove or change the configuration, they will appear with a warning icon. Platform Settings that will no longer work on your device are marked with a caution icon and require your acknowledgement.



Step 4 If you need to manually remove or change any listed configurations, do the following.

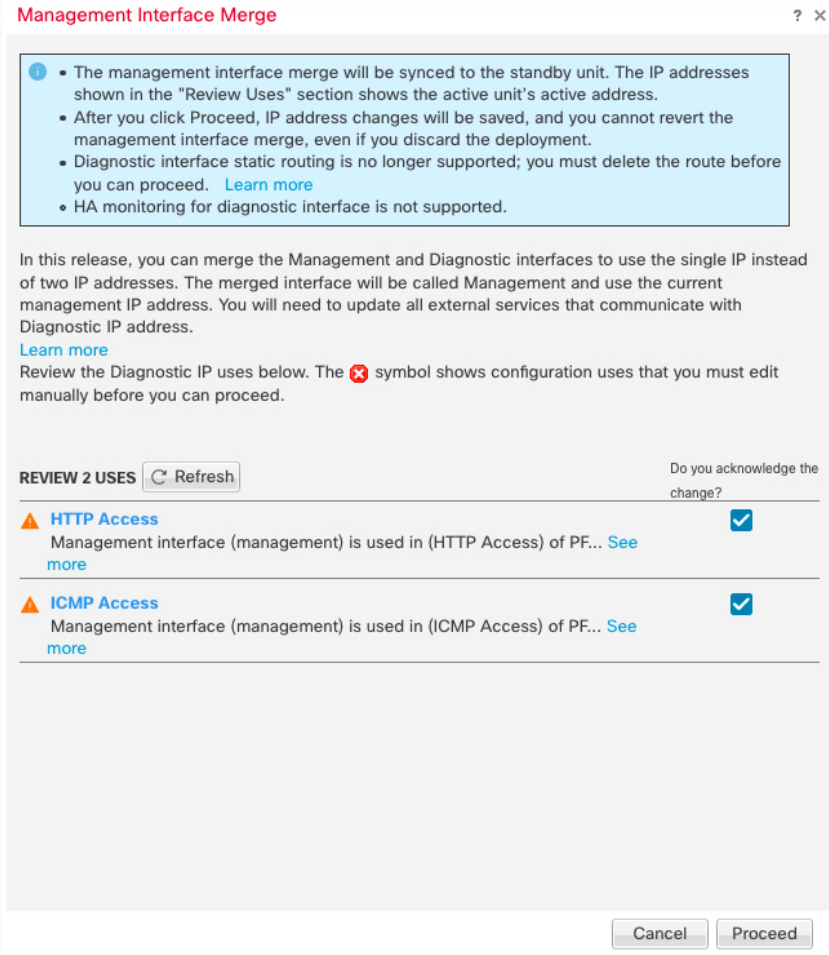
a) Click **Cancel** to close the **Management Interface Merge** dialog box.

b) Navigate to the feature area. You can then delete the item, or choose a data interface instead.

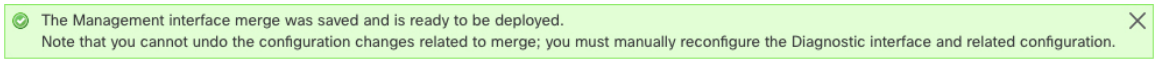
c) Reopen the **Management Interface Merge** dialog box.

There should no longer be any warnings.

Step 5 For each configuration caution, click the box in **Do you acknowledge the change?** column, and then click **Proceed**.



After the configuration is merged, you see a success banner:



Step 6 Deploy the new merged configuration.

Caution
After you deploy the merged configuration, you can unmerge the interfaces from Cloud-Delivered Firewall Management Center; however the Diagnostic interface will have to be reconfigured manually. See Unmerge the management interface, on page 1034. Also, if you restore a configuration that is unmerged, or roll back to an unmerged configuration, then the device will revert to that unmerged configuration.

After the merge, the Management interface is shown on the **Interfaces** page, although it is read-only.

Step 7 After the merge, if you had any external services that communicated with the Diagnostic interface, you need to change their configuration to use the Management interface IP address.

For example:

- SNMP client

- RADIUS server—RADIUS servers often verify the IP address for incoming traffic, so you need to change that IP address to the Management address. Moreover, for a High Availability pair, you need to allow both the primary and secondary Management IP addresses; the Diagnostic interface used to support a single "floating" IP address that stayed with the active unit, but Management does not support that functionality.

Unmerge the management interface

This task separates the merged Management and Diagnostic interface into two distinct interfaces. This is useful when migrating network configurations or when separate interface management is required for specific network architectures.

The Firewall Threat Defense 7.4 and later supports a merged Management and Diagnostic interface. If you need to unmerge your interfaces, perform this procedure. We recommend using unmerged mode temporarily while you migrate your network to a merged mode deployment. Separate Management and Diagnostic interfaces may not be supported in all future releases.

Unmerging the interfaces does not restore your original Diagnostic configuration (if you upgraded and then merged your interfaces). You will need to reconfigure the Diagnostic interface manually. Also, the Management interface will now be named "management"; you cannot rename it "diagnostic."

Alternatively, if you used the Backup function to save an old unmerged configuration, you can restore that configuration or you can use the or Cloud-Delivered Firewall Management Center configuration rollback feature, and the device will be in an unmerged state with the Diagnostic configuration intact.

Before you begin

- To view the current mode of the device, enter the **show management-interface convergence** command at the Firewall Threat Defense CLI. The following output shows that the Management interfaces are merged:

```
> show management-interface convergence
management-interface convergence
>
```

  The following output shows that the Management interfaces are not merged:

```
> show management-interface convergence
no management-interface convergence
>
```

- For High Availability pairs and clusters, perform this task on the active/control unit. The merged configuration will be replicated automatically to the standby/data units.

Follow these steps to unmerge the Management Interface:

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) for your Firewall Threat Defense. The **Interfaces** page is selected by default.

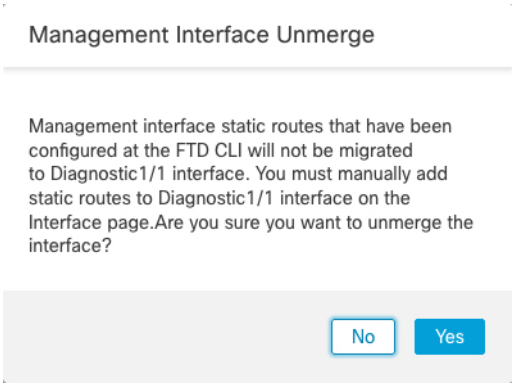
Step 2 For the Management interface, click **Unmerge Management Interface** ( ).

Figure 241: Management interface selection



Step 3 Click **Yes** to confirm that you want to unmerge the interface.

Figure 242: Unmerge confirmation



Step 4 Deploy the new unmerged configuration.

Note
If you restore a configuration that is merged, or roll back to a merged configuration, then the device will revert to that merged configuration.

After the merge, the Management interface is no longer shown on the **Interfaces** page.

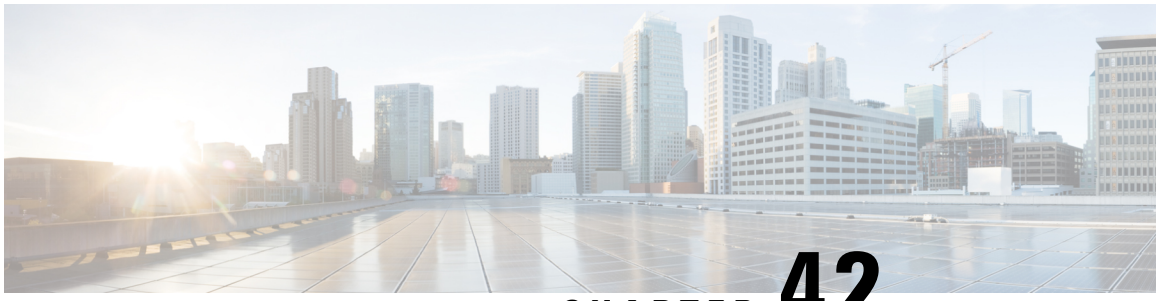
History for interfaces

This reference provides the complete version history for interface-related features, enhancements, and changes across different versions of the Secure Firewall system.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Sync Device is now called **Sync Interfaces**	20250219	7.7.0	**Sync Device** was changed to **Sync Interfaces** to indicate that this function is only for interface changes. This function no longer detects changes made to the manager access interface; see **Devices** > **Device Management** > **Device** > **Management** > **Manager Access Details: Configuration**. Other out-of-band configuration changes performed at the diagnostic CLI in recovery-config mode need to be discovered at **Devices** > **Device Management** > **Device** > **Health** > **Out of Band Status**. New/Modified screens: **Devices** > **Device Management** > **Interfaces**
Loopback and Management type interface group objects	7.4.0	7.4.0	You can now create interface group objects that include only management-only interfaces or only loopback interfaces. You can then use these groups for management features such as DNS servers, HTTP access, or SSH. Loopback groups are supported for any feature that supports loopback interfaces. Note that DNS does not support management interfaces. New/Modified screens: **Objects** > **Object Management** > **Interface** > **Add** > **Interface Group**
Merged Management and Diagnostic interfaces	7.4.0	7.4.0	For new devices using 7.4 and later, you cannot use the legacy Diagnostic interface. Only the merged Management interface is available. If you upgraded to 7.4 or later, and you did not have any configuration for the Diagnostic interface, then the interfaces will merge automatically. If you upgraded to 7.4 or later, and you have configuration for the Diagnostic interface, then you have the choice to merge the interfaces manually, or you can continue to use the separate Diagnostic interface. Note that support for the Diagnostic interface will be removed in a later release, so you should plan to merge the interfaces as soon as possible. Merged mode also changes the behavior of AAA traffic to use the data routing table by default. The management-only routing table can now only be used if you specify the management-only interface (including Management) in the configuration. New/Modified screens: **Devices** > **Device Management** > **Interfaces** New/Modified commands: **show management-interface convergence**
Default Forward Error Correction (FEC) on Secure Firewall 3100 fixed ports changed to Clause 108 RS-FEC from Clause 74 FC-FEC for 25 GB+ SR, CSR, and LR transceivers	7.2.4	7.2.4	When you set the FEC to Auto on the Secure Firewall 3100 fixed ports, the default type is now set to Clause 108 RS-FEC instead of Clause 74 FC-FEC for 25 GB+ SR, CSR, and LR transceivers.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
LLDP support for the Firepower 2100, Secure Firewall 3100	7.2.0	7.2.0	You can enable Link Layer Discovery Protocol (LLDP) for Firepower 2100 and Secure Firewall 3100 interfaces. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **Network Connectivity** New/Modified commands: **show LLDP status**, **show LLDP neighbors**, **show LLDP statistics**
Pause Frames for Flow Control for the Secure Firewall 3100	7.2.0	7.2.0	If you have a traffic burst, dropped packets can occur if the burst exceeds the buffering capacity of the FIFO buffer on the NIC and the receive ring buffers. Enabling pause frames for flow control can alleviate this issue. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **Network Connectivity**
Support for hot swapping the network module for the Secure Firewall 3100	7.1.0	7.1.0	You can add or remove the network module on the Secure Firewall 3100 while the firewall is powered up. To replace a module with another module of the same type, you do not need to reboot. After initial bootup, adding a module, permanently removing a module, or replacing a module with a new type requires a reboot. New/Modified screens: **Devices** > **Device Management** > **Chassis Operations**
Support for Forward Error Correction for the Secure Firewall 3100	7.1.0	7.1.0	Secure Firewall 3100 25 Gbps interfaces support Forward Error Correction (FEC). FEC is enabled by default and set to Auto. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** > **Hardware Configuration**
Support for setting the speed based on the SFP for the Secure Firewall 3100	7.1.0	7.1.0	The Secure Firewall 3100 supports speed detection for interfaces based on the SFP installed. Detect SFP is enabled by default. This option is useful if you later change the network module to a different model, and want the speed to update automatically. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** > **Hardware Configuration**
LLDP support for the Firepower 1100.	7.1.0	7.1.0	You can enable Link Layer Discovery Protocol (LLDP) for Firepower 1100 interfaces. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **LLDP** New/Modified commands: **show LLDP status**, **show LLDP neighbors**, **show LLDP statistics**

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Interface auto-negotiation is now set independently from speed and duplex, interface sync improved.	7.1.0	7.1.0	Interface auto-negotiation is now set independently from speed and duplex. Also, when you sync the interfaces in Cloud-Delivered Firewall Management Center, hardware changes are detected more effectively. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **Speed** Supported platforms: Firepower 1000, 2100, Secure Firewall 3100
Firepower 1100/2100 series fiber interfaces now support disabling auto-negotiation.	6.7.0	6.7.0	You can now configure a Firepower 1100/2100 series fiber interface to disable flow control and link status negotiation. Previously, when you set the fiber interface speed (1000 or 10000 Mbps) on these devices, flow control and link status negotiation was automatically enabled. You could not disable it. Now, you can deselect **Auto-negotiation** and set the speed to 1000 to disable flow control and link status negotiation. You cannot disable negotiation at 10000 Mbps. New/modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **Speed**



CHAPTER 42

Regular Firewall Interfaces

This chapter includes regular firewall Firewall Threat Defense interface configuration including EtherChannels, VLAN subinterfaces, IP addressing, and more.

- Requirements and prerequisites for regular firewall interfaces, on page 1039
- Configure switch ports, on page 1040
- Loopback interfaces, on page 1052
- Guidelines for Loopback Interfaces, on page 1053
- Configure a loopback interface, on page 1053
- Rate-Limit traffic to the loopback interface, on page 1054
- VLAN subinterfaces and 802.1Q trunking, on page 1058
- Configure VXLAN Interfaces, on page 1062
- Configure routed and transparent mode interfaces, on page 1077
- Advanced interface settings, on page 1103
- History for Regular Firewall Interfaces, on page 1115

Requirements and prerequisites for regular firewall interfaces

This reference provides the requirements and prerequisites for configuring regular firewall interfaces, including supported models and required user roles.

Model support

Firewall Threat Defense

User roles

- Admin
- Access Admin
- Network Admin

Configure switch ports

For supported platforms, you can configure each interface to run as a regular firewall interface or as a Layer 2 hardware switch port. This section includes tasks for starting your switch port configuration, including enabling or disabling the switch mode and creating VLAN interfaces and assigning them to switch ports. This section also describes how to customize Power over Ethernet (PoE) on supported interfaces.

About switch ports

This section describes the switch ports for models with Layer 2 switches.

Switch ports and interfaces

Switch ports and firewall operate at different OSI layers to handle traffic forwarding.

Port and interface types

For each physical interface, you can set its operation as a firewall interface or as a switch port. This information describes physical interface and port types as well as logical VLAN interfaces to which you assign switch ports:

- Physical firewall interface—In routed mode, these interfaces forward traffic between networks at Layer 3, using the configured security policy to apply firewall and VPN services. In transparent mode, these interfaces are bridge group members that forward traffic between the interfaces on the same network at Layer 2, using the configured security policy to apply firewall services. In routed mode, you can also use Integrated Routing and Bridging with some interfaces as bridge group members and others as Layer 3 interfaces. By default, the Ethernet 1/1 interface is configured as a firewall interface. You can also configure these interfaces to be IPS-only (passive interfaces).
- Physical switch port—Switch ports forward traffic at Layer 2, using the switching function in hardware. Switch ports on the same VLAN can communicate with each other using hardware switching, and traffic is not subject to the Firewall Threat Defense security policy. Access ports accept only untagged traffic, and you can assign them to a single VLAN. Trunk ports accept untagged and tagged traffic, and can belong to more than one VLAN. By default, Ethernet 1/2 and subsequent ports are configured as access switch ports on VLAN 1. You cannot configure the Management interface as a switch port.
- Logical VLAN interface—These interfaces operate the same as physical firewall interfaces, with the exception that you cannot create subinterfaces, IPS-only interfaces (inline sets and passive interfaces), or EtherChannel interfaces. When a switch port needs to communicate with another network, then the Firewall Threat Defense device applies the security policy to the VLAN interface and routes to another logical VLAN interface or firewall interface. You can even use Integrated Routing and Bridging with VLAN interfaces as bridge group members. Traffic between switch ports on the same VLAN are not subject to the Firewall Threat Defense security policy, but traffic between VLANs in a bridge group are subject to the security policy, so you may choose to layer bridge groups and switch ports to enforce the security policy between certain segments.

Power Over Ethernet

PoE is available on these ports:

- Secure Firewall 1210CP—Ethernet 1/5, 1/6, 1/7, and 1/8 using IEEE 802.3af (PoE), 802.3at (PoE+) up to 30 watts per port, up to a combined 120 watts.

PoE+ uses Link Layer Discovery Protocol (LLDP) to negotiate the power level. Power is only supplied when needed.

If you shut down the interface, then you disable power to the device.

Auto-MDI/MDIX feature

Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase.

Auto-MDI/MDIX configuration details

For all switch ports, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. The interface enables Auto-MDI/MDIX when either the speed or duplex is set to auto-negotiate. If you set both the speed and duplex to fixed values, auto-negotiation is disabled, and Auto-MDI/MDIX is also disabled. When the speed and duplex are set to 1000 and full, the interface always auto-negotiates. Auto-MDI/MDIX enabled in this state, and disabling it is not possible.

Prerequisites for switch ports

This reference provides model support and other prerequisites for switch ports.

Model support

- Secure Firewall 200
- Firepower 1010
- Secure Firewall 1210/1220

Guidelines for switch ports

Context mode restrictions

Multiple context mode is not supported on several firewall models, with specific exceptions for the Secure Firewall 1210/1220.

- The Secure Firewall 220 does not support multiple context mode.
- The Firepower 1010 does not support multiple context mode.
- The Secure Firewall 1210/1220 does not support multiple context mode.

High availability and clustering limitations

Do not use switch port functionality when using High availability because switch ports operate in hardware and continue to pass traffic on both active and standby units.

- No cluster support.

- Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. High availability is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal High availability network setup, active switch ports on both units will lead to network loops. We suggest that you use external switches for any switching capability. Note that VLAN interfaces can be monitored by failover, while switch ports cannot. Theoretically, you can put a single switch port on a VLAN and successfully use High availability, but a simpler setup is to use physical firewall interfaces instead.
- You can only use a firewall interface as the failover link.

Logical VLAN interfaces (SVIs) configuration requirements

Ensure unique VLAN IDs between logical VLAN interfaces and VLAN subinterfaces on firewall interfaces to prevent conflicts.

- If you also use VLAN subinterfaces on a firewall interface, you cannot use the same VLAN ID as for a logical VLAN interface.
- MAC Addresses:
  - Routed firewall mode—All VLAN interfaces share a MAC address. Ensure that any connected switches can support this scenario. If the connected switches require unique MAC addresses, you can manually assign MAC addresses. See Configure the MAC address, on page 1110.
  - Transparent firewall mode—Each VLAN interface has a unique MAC address. You can override the generated MAC addresses if desired by manually assigning MAC addresses. See Configure the MAC address, on page 1110.

Bridge groups restrictions

You cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.

VLAN interface and switch port unsupported features

VLAN interfaces and switch ports do not support these features.

- Dynamic routing
- Multicast routing
- Equal-Cost Multi-Path routing (ECMP)
- Inline sets or Passive interfaces
- EtherChannels—Switch ports cannot be part of an EtherChannel. PoE is also not supported on a port in an EtherChannel.
- Failover and state link
- Security group tagging (SGT)

Other guidelines and limitations

Configure a maximum of 60 named interfaces and ensure the Management interface is not configured as a switch port.

Default settings

Know the default interface assignments for each firewall model to properly plan your configuration.

- Ethernet 1/1 is a firewall interface.
- On 1010, Ethernet 1/2 through Ethernet 1/8 are switch ports assigned to VLAN 1.
- On 1210, Ethernet 1/2 through Ethernet 1/8 are switch ports assigned to VLAN 1.
- On 1220, Ethernet 1/2 through Ethernet 1/10 are switch ports assigned to VLAN 1.
- On 220, Ethernet 1/2 through Ethernet 1/5 are switch ports assigned to VLAN 1.
- Default Speed and Duplex—By default, the speed and duplex are set to auto-negotiate.

Configure switch ports and power over ethernet

To configure switch ports and PoE, complete the following tasks.

Enable or disable switch port mode

Set each interface independently to be either a firewall interface or a switch port based on your network requirements.

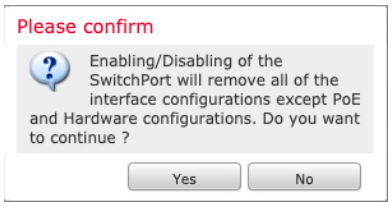
You can set each interface independently to be either a firewall interface or a switch port. By default, Ethernet 1/1 is a firewall interface, and the remaining Ethernet interfaces are configured as switch ports.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Set the switch port mode by clicking the slider in the **SwitchPort** column so it shows as **Slider enabled** () or **Slider disabled** ().

By default, switch ports are set to access mode in VLAN 1. You must manually add a logical VLAN 1 interface (or whichever VLAN you set for these switch ports) for traffic to be routed and to participate in the Firewall Threat Defense security policy (see Configure a VLAN interface, on page 1044). You cannot set the Management interface to switch port mode. When you change the switch port mode, all unsupported configuration is removed:



In addition, the interface is disabled if it was already enabled. Be sure to reenable the interface.

Configure a VLAN interface

Configure VLAN interfaces to enable traffic routing and participation in the Firewall Threat Defense security policy.

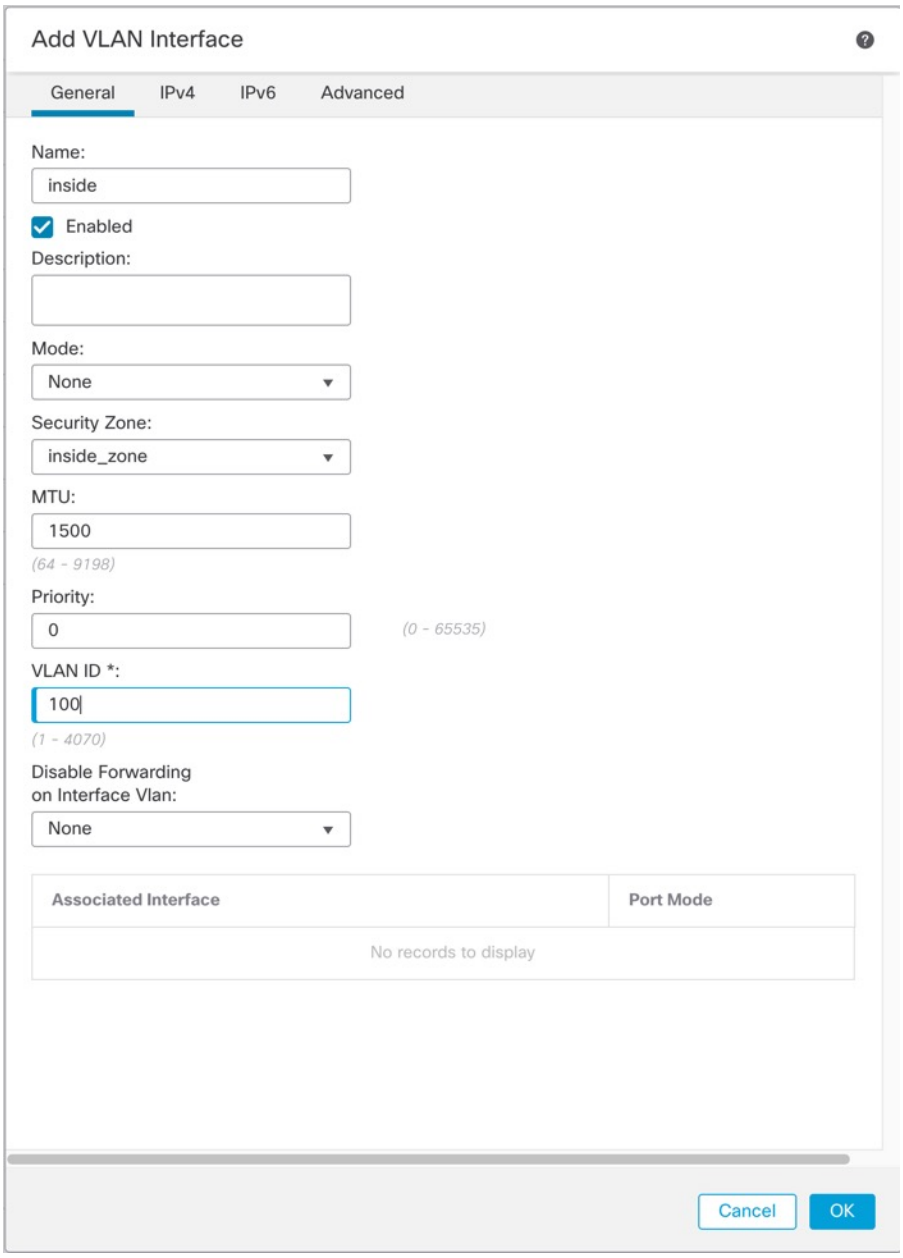
By default, switch ports are assigned to VLAN1; however, you must manually add the logical VLAN1 interface (or whichever VLAN you set for these switch ports) for traffic to be routed and to participate in the Firewall Threat Defense security policy.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Add Interfaces** > **VLAN Interface**.

Step 3 On **General**, set the following VLAN-specific parameters:



If you are editing an existing VLAN interface, the **Associated Interface** table shows switch ports on this VLAN.

a) Set the **VLAN ID** between 1 and 4070, excluding IDs in the range 3968 to 4047, which are reserved for internal use.

   You cannot change the VLAN ID after you save the interface; the VLAN ID is both the VLAN tag used, and the interface ID in your configuration.

b) (Optional) Choose a VLAN ID for **Disable Forwarding on Interface VLAN** to disable forwarding to another VLAN.

For example, you have one VLAN assigned to the outside for internet access, one VLAN assigned to an inside business network, and a third VLAN assigned to your home network. The home network does not need to access the business network, so you can disable forwarding on the home VLAN. The business network can access the home network, but the home network cannot access the business network.

Step 4 To complete the interface configuration, see one of these procedures:

- Configure routed mode interfaces, on page 1080
- Configure general bridge group member interface parameters, on page 1086

Step 5 Click **OK**.

Step 6 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure switch ports as access ports

Configure switch ports as access ports to assign them to a single VLAN for network segmentation and traffic management.

To assign a switch port to a single VLAN, configure it as an access port. Access ports accept only untagged traffic. The interfaces that are enabled with switch ports and assigned to VLAN1 are as follows:

Device Models	Switch Port Interfaces
Secure Firewall 220	Ethernet 1/2 through Ethernet 1/5
Firepower 1010	Ethernet 1/2 through Ethernet 1/8
Secure Firewall 1210	Ethernet 1/2 through Ethernet 1/8
Secure Firewall 1220	Ethernet 1/2 through Ethernet 1/10

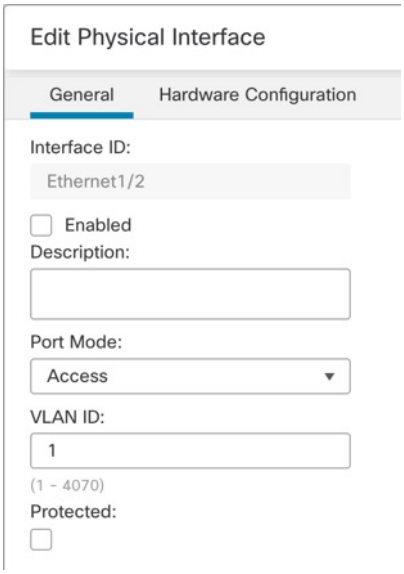
Note The device does not support Spanning Tree Protocol for loop detection in the network. Therefore you must ensure that any connection with the Firewall Threat Defense does not end up in a network loop.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Figure 243: Edit physical interface



Step 3 Enable the interface by checking the **Enabled** check box.

Step 4 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 5 Set the **Port Mode** to **Access**.

Step 6 In the **VLAN ID** field, set the VLAN for this switch port, between 1 and 4070.

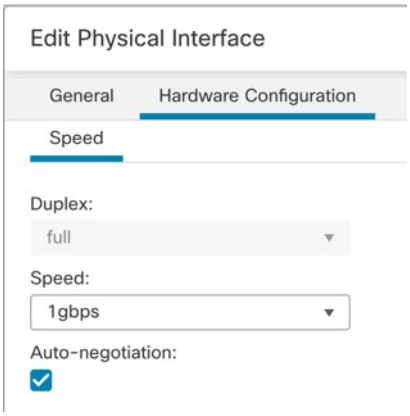
The default VLAN ID is 1.

Step 7 (Optional) Check the **Protected** check box to set this switch port as protected, so you can prevent the switch port from communicating with other protected switch ports on the same VLAN.

You might want to prevent switch ports from communicating with each other if: the devices on those switch ports are primarily accessed from other VLANs; you do not need to allow intra-VLAN access; and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you enable **Protected** on each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.

Step 8 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

Figure 244: Hardware configuration



Check the **Auto-negotiation** check box (the default) to auto-detect the speed and duplex. If you uncheck it, you can set the speed and duplex manually:

- **Duplex**—Choose **Full** or **Half**.
- **Speed**—Choose **10mbps**, **100mbps**, or **1gbps**.

Note
If a peer switch connecting to the port over a 50G cable does not support auto-negotiation, ensure to disable auto-negotiation on the switch and the Threat Defense interface as well. For example, N9K-C93400LD-H1 does not support auto-negotiation on a 50G cable. Hence, you must disable the default auto-negotiation on the platform and the switch for the port to be connected.

Step 9 Click **OK**.

Click **Save**.

What to do next

Go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure switch ports as trunk ports

This task creates a trunk port that can carry multiple VLANs using 802.1Q tagging.

This procedure describes how to create a trunk port that can carry multiple VLANs using 802.1Q tagging. Trunk ports accept untagged and tagged traffic. Traffic on allowed VLANs pass through the trunk port unchanged.

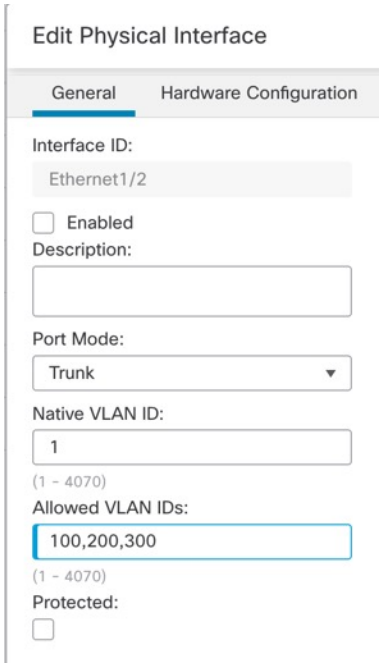
When the trunk receives untagged traffic, it tags it to the native VLAN ID so that the Firewall Threat Defense device can forward the traffic to the correct switch ports or can route it to another firewall interface. When the Firewall Threat Defense device sends native VLAN ID traffic out of the trunk port, it removes the VLAN tag. Be sure to set the same native VLAN on the trunk port on the other switch so that the untagged traffic will be tagged to the same VLAN.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Figure 245: Set trunk port mode



Step 3 Enable the interface by checking the **Enabled** check box.

Step 4 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 5 Set the **Port Mode** to **Trunk**.

Step 6 Set the VLAN IDs:

a) In the **Native VLAN ID** field, set the native VLAN for this switch port, between 1 and 4070.

The default native VLAN ID is 1.

Each port can only have one native VLAN, but every port can have either the same or a different native VLAN.

b) In the **Allowed VLAN IDs** field, enter the VLANs for this trunk port between 1 and 4070.

You can identify up to 20 IDs in one of these ways:

- A single number (n)
- A range (n-x)

- Numbers and ranges separated by commas, for example:

  5,7-10,13,45-100

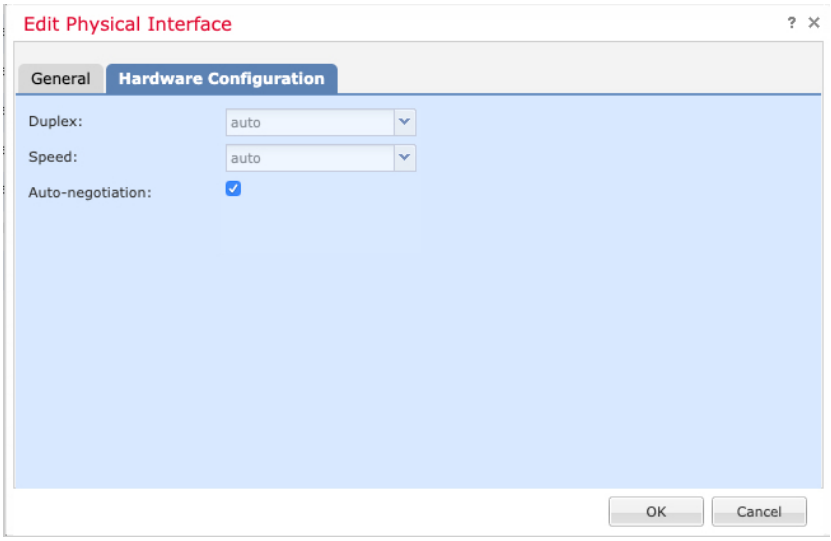
  You can enter spaces instead of commas.

If you include the native VLAN in this field, it is ignored; the trunk port always removes the VLAN tagging when sending native VLAN traffic out of the port. Moreover, it will not receive traffic that still has native VLAN tagging.

Step 7 (Optional) Check the **Protected** check box to set this switch port as protected, so you can prevent the switch port from communicating with other protected switch ports on the same VLAN.

You might want to prevent switch ports from communicating with each other if: the devices on those switch ports are primarily accessed from other VLANs; you do not need to allow intra-VLAN access; and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you enable **Protected** on each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.

Step 8 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.



Check the **Auto-negotiation** check box (the default) to auto-detect the speed and duplex. If you uncheck it, you can set the speed and duplex manually:

- **Duplex**—Choose **Full** or **Half**.
- **Speed**—Choose **10mbps**, **100mbps**, or **1gbps**.

If a peer switch connecting to the port over a 50G cable does not support auto-negotiation, ensure to disable auto-negotiation on the switch and the Threat Defense interface as well. For example, N9K-C93400LD-H1 does not support auto-negotiation on a 50G cable. Hence, you must disable the default auto-negotiation on the platform and the switch for the port to be connected.

Step 9 Click **OK**.

Click **Save**.

What to do next

Go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Power over Ethernet

Configure Power over Ethernet (PoE) settings to control power delivery to connected devices and optimize power allocation.

Power over Ethernet (PoE) ports provide power for devices such as IP phones or wireless access points. PoE is enabled by default. This procedure describes how to disable and enable PoE and how to set optional parameters.

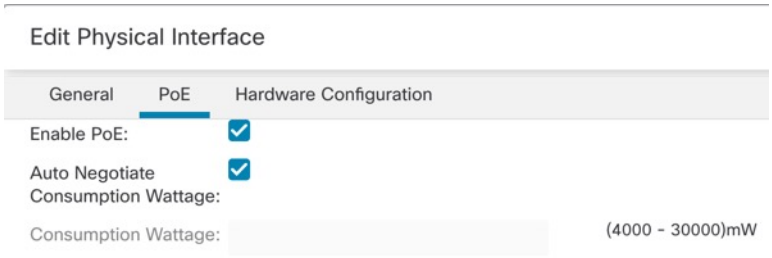
Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for Ethernet1/7 or 1/8 on Firepower 1010 or for any interface from Ethernet 1/5-1/8 on Secure Firewall 1210CP .

Step 3 Click **PoE**.

Figure 246: PoE



Step 4 Check the **Enable PoE** check box.

PoE is enabled by default.

Step 5 Choose autonegotiation or manual power.

- **Auto Negotiate Consumption Wattage**—Delivers power automatically to the powered device using a wattage appropriate to the class of the powered device. The firewall uses LLDP to further negotiate the correct wattage. When you connect a device of a certain class, it will be provisioned up to the maximum for that class in case it ever needs to use more power. For example, if you add a class 4 device that requests 12.95W, it will be allocated 30W even if it doesn't currently use all that power. Some devices can renegotiate power needs. If you know your device needs less power than is allocated, you can instead set the **Consumption Wattage** manually to free up the power for other devices.

- **Consumption Wattage**—Uncheck the **Auto Negotiate Consumption Wattage** check box to set the manual wattage to manually specify the wattage in milliwatts, from 4000 to 30000. Use this option if you want to set the watts manually and disable LLDP negotiation. For manual allocation, the class will show in **show power inline** output as *n/a* because the class is not used to decide the power consumption.

  View the current PoE status using the **show power inline** command.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Loopback interfaces

A loopback interface is a software-only interface that emulates a physical interface. This interface is reachable on IPv4 and IPv6 through multiple physical interfaces. The loopback interface helps to overcome path failures. It is accessible from any physical interface. If one interface goes down, you can access the loopback interface from another.

Loopback interfaces can be used for:

- AAA
- BGP
- DNS
- HTTP
- ICMP
- IPsec flow offload—Secure Firewall 1200/3100/4200 only
- NetFlow
- SNMP
- SSH
- Static and dynamic VTI tunnels
- Syslog

The Firewall Threat Defense can distribute the loopback address using dynamic routing protocols, or you can configure a static route on the peer device to reach the loopback IP address through one of the Firewall Threat Defense's physical interfaces. You cannot configure a static route on the Firewall Threat Defense that specifies the loopback interface.

Related Topics

Guidelines for Loopback Interfaces, on page 1053
Configure a loopback interface, on page 1053

Guidelines for Loopback Interfaces

This topic outlines the specific constraints and supported configurations for loopback interfaces in various network environments.

Firewall Mode

- Supported in routed mode only.

High availability and Clustering

- No clustering support.

Additional Guidelines and Limitations

- TCP sequence randomization is always disabled for traffic from the physical interface to the loopback interface.

Configure a loopback interface

Configure a loopback interface to provide a virtual interface that remains active regardless of physical interface status, commonly used for routing protocols and management purposes.

Loopback interfaces are virtual interfaces that provide stable IP addresses for device management and routing protocols. They remain active as long as the device is operational, making them ideal for network management tasks.

Before you begin

Follow these steps to add a loopback interface for a device:

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 From the **Add Interfaces** drop-down list, choose **Loopback Interface**.

Step 3 In the **General** tab, configure the following parameters:

a) **Name**—Enter a name for the loopback interface.
b) **Enabled**—Check the check box to enable the loopback interface.
c) **Loopback ID**—Enter the loopback ID between 1 to 1024.
d) **Description**—Enter a description for the loopback interface.

Step 4 Configure the routed mode interface parameters. See Configure routed mode interfaces, on page 1080.

Rate-Limit traffic to the loopback interface

You should rate-limit traffic going to the loopback interface IP address to prevent excessive load on the system. You can add a connection limit rule to the global service policy.

Rate-limiting protects the system from Denial of Service (DoS) attacks and ensures stable performance by controlling the number of connections to critical interface addresses.

Before you begin

Follow these steps to rate-limit traffic to the loopback interface:

Procedure

Step 1 Create an extended access list identifying traffic to the loopback interface IP address(es).

a. Choose **Objects** > **Object Management** > **Access List** > **Extended**.

b. Click **Add Extended Access List** to create a new ACL.

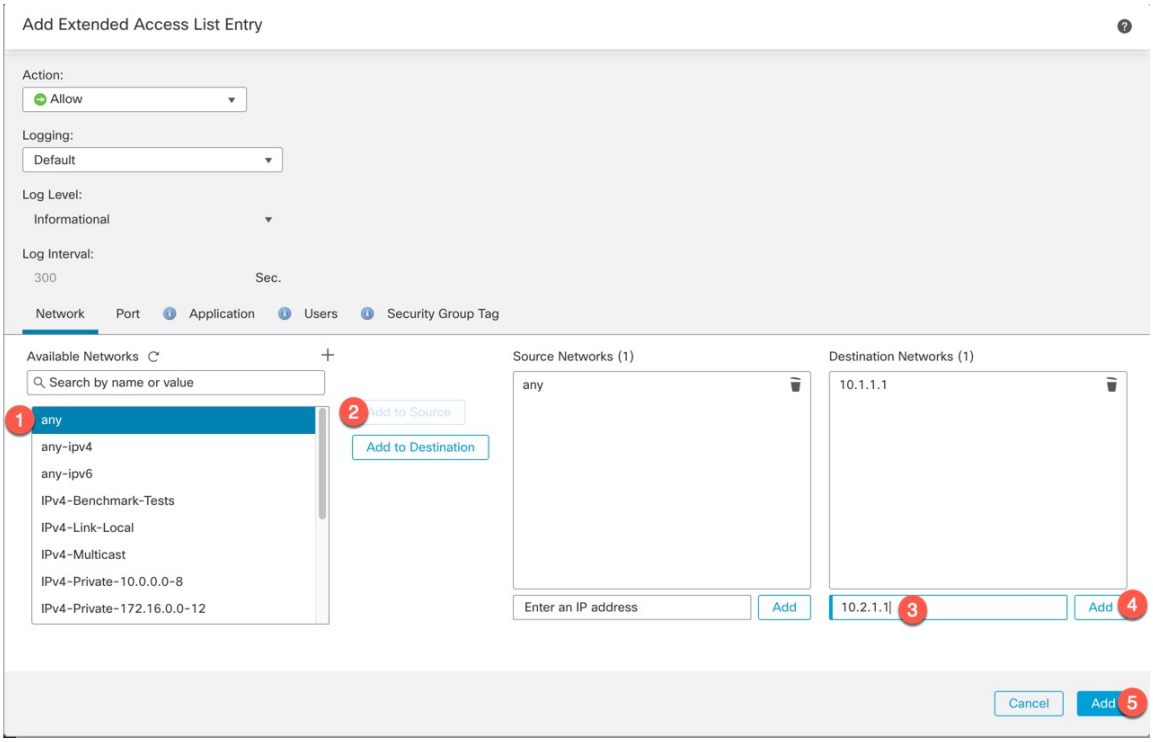
c. In the **New Extended Access List Object** dialog box, enter a name for the ACL (no spaces allowed), and click **Add** to create a new entry.

Figure 247: Name ACL and add entry



d. Configure the source (any) and destination addresses (loopback IP addresses) on the **Network** tab.

Figure 248: Source and destination networks



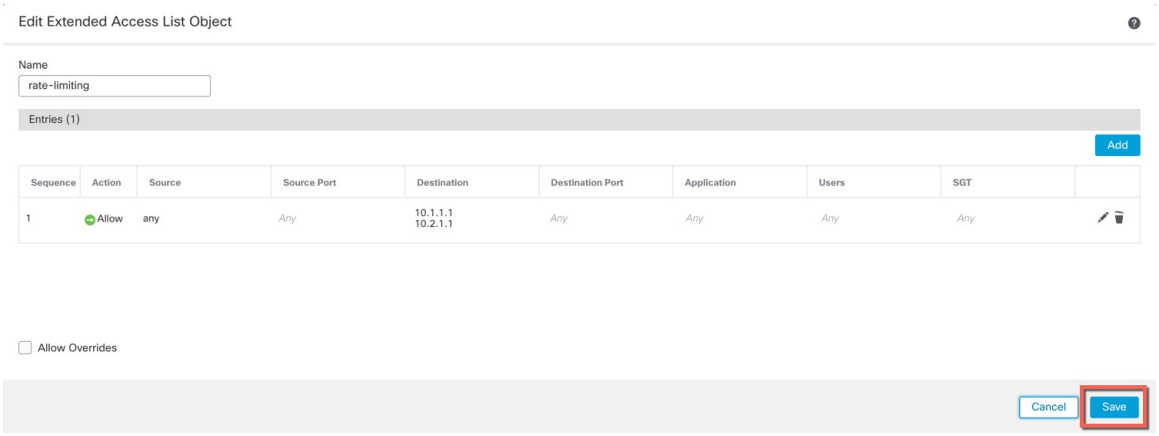
Note
Keep the default **Action** as **Allow** (match) and other settings as-is.

- Source—Select **any** from the **Available Networks** list, and click **Add to Source**. You can also narrow this access list by specifying the source IP addresses instead of **any**.
- Destination—Type an address in the edit box below the **Destination Networks** list and click **Add**. Repeat for each loopback interface.

e. Click **Add** to add the entry to the ACL.

f. Click **Save** to save the ACL.

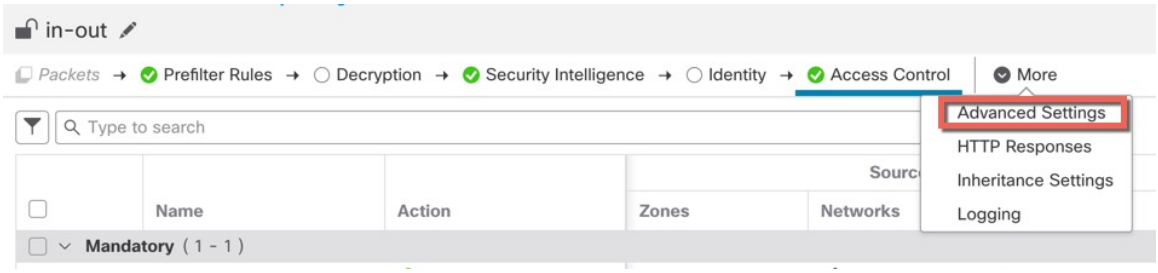
Figure 249: Save ACL



Step 2 Choose **Policies** > **Access Control heading** > **Access Control**, and click **Edit** (✎) for the access control policy assigned to your device.

Step 3 Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Figure 250: Advanced settings



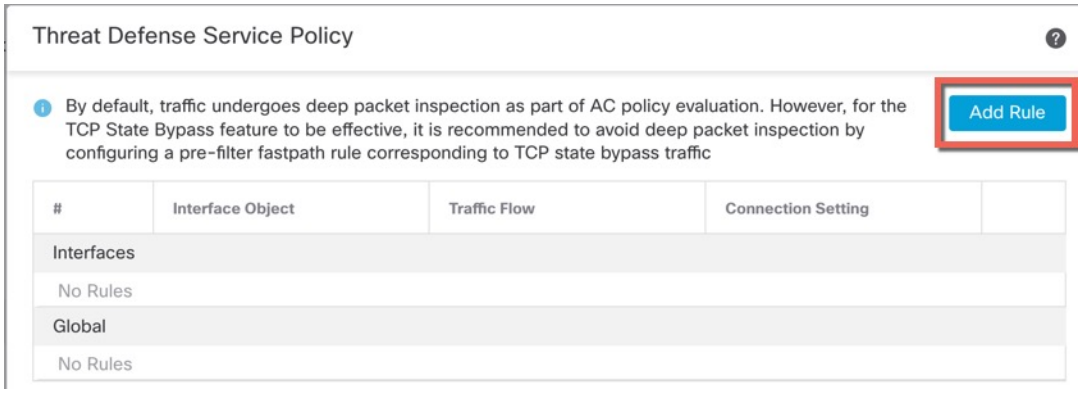
Step 4 Click **Edit** (✎) in the **Threat Defense Service Policy** group.

Figure 251: Threat defense service policy



Step 5 Click **Add Rule** to create a new rule.

Figure 252: Add rule



The service policy rule wizard opens to step you through the process of configuring the rule.

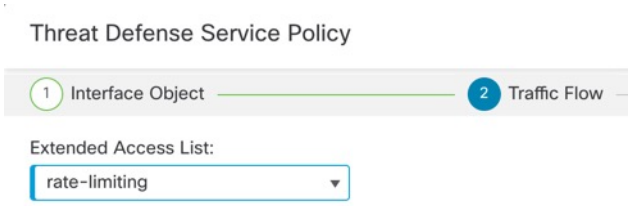
Step 6 In the **Interface Object** step, click **Global** to create a global rule, which applies to all interfaces, then click **Next**.

Figure 253: Global policy



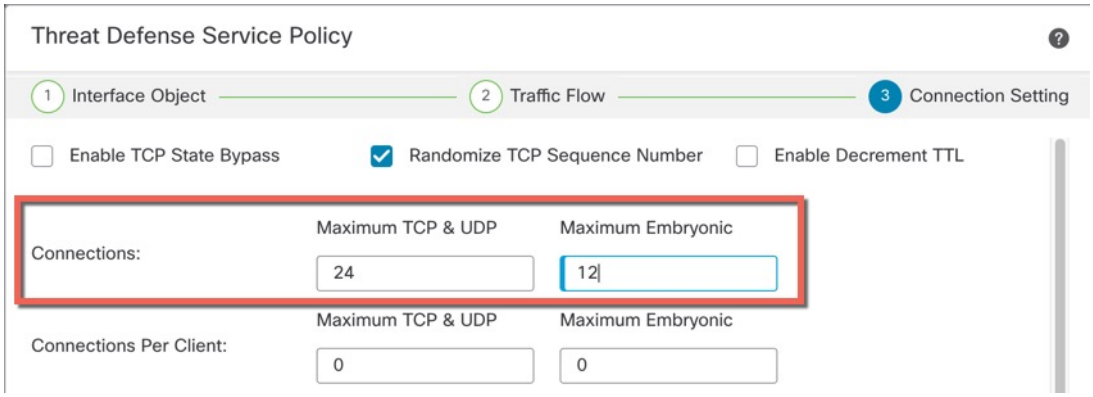
Step 7 In the **Traffic Flow** step, select the extended access list object you created in Step 1, on page 1054, and then click **Next**.

Figure 254: Choose extended access list



Step 8 In the **Connection Setting** step, set the **Connections** limits.

Figure 255: Set connection limits



Set the **Maximum TCP & UDP** connections to the expected number of connections for the loopback interface, and the **Maximum Embryonic** connections to a lower number. For example, you can set it to 5/2, or 10/5, or 1024/512, depending on the expected loopback interface sessions you need.

Setting the embryonic connection limit enables TCP Intercept, which protects the system from a DoS attack perpetrated by flooding an interface with TCP SYN packets.

Step 9 Click **Finish** to save your changes.

Click **OK**.

Click **Save** on the **Advanced Settings** window.

What to do next

Deploy the changes to the affected devices.

VLAN subinterfaces and 802.1Q trunking

VLAN subinterfaces let you divide a physical, redundant, or EtherChannel interface into multiple logical interfaces that are tagged with different VLAN IDs. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk. Because VLANs let you keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or devices.

Guidelines for VLAN subinterfaces

Model support

- 1010/200/1210/1220—VLAN subinterfaces are not supported on switch ports or VLAN interfaces.

High availability and clustering

Do not use a subinterface for the failover or state link or for the cluster control link. The exception is for multi-instance mode: you can use a *chassis*-defined subinterface for these links.

Additional guidelines

- Preventing untagged packets on the physical interface—If you use subinterfaces, disable the physical interface from passing traffic. The physical interface transmits untagged packets. This behavior also applies to the active physical interface in a redundant interface pair and also to EtherChannel links. For the subinterface to pass traffic, the physical, redundant, or EtherChannel interface must be enabled. If you want to let the physical, redundant, or EtherChannel interface pass untagged packets, you can configure the name as usual.
- Configure subinterfaces only on non-Management interfaces. Avoid using the dedicated Management interface configured at the CLI or a data interface used for manager access.
- All subinterfaces on the same parent interface must be either members of bridge groups or routed interfaces; you cannot mix and match.
- The Firewall Threat Defense does not support the Dynamic Trunking Protocol (DTP), so configure the connected switch port to trunk unconditionally.
- Assign unique MAC addresses to the subinterfaces that are defined on the Firewall Threat Defense. Subinterfaces use the same burned-in MAC address of the parent interface by default. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the Firewall Threat Defense.



Note If you manually assign a MAC address, be sure to assign MAC addresses to all subinterfaces on the same physical interface to avoid unexpected behavior and outages.

Maximum number of VLAN subinterfaces by device model

Your device model limits the maximum number of VLAN subinterfaces that you can configure. Note that you can configure subinterfaces on data interfaces only, you cannot configure them on the management interface.

This table explains the limits for each device model.

Model	Maximum VLAN Subinterfaces
Secure Firewall 200	1024
Firepower 1010	60
Firepower 1120	512
Firepower 1140, 1150	1024
Secure Firewall 1200	1024
Secure Firewall 3100	1024

Model	Maximum VLAN Subinterfaces
Secure Firewall 4200	1024
Firepower 4100	1024
Firepower 9300	1024
Firewall Threat Defense Virtual	50
ISA 3000	100

Add a subinterface

This task enables you to add subinterfaces to physical, redundant, or port-channel interfaces for network segmentation using VLAN tagging.

Add one or more subinterfaces to a physical, redundant, or port-channel interface.

For the Firepower 4100/9300, you can configure subinterfaces in FXOS for use with container instances; see Add a VLAN subinterface for container instances. These subinterfaces appear in the Cloud-Delivered Firewall Management Center interface list. You can also add subinterfaces in Cloud-Delivered Firewall Management Center, but only on parent interfaces that do not already have subinterfaces defined in FXOS.



Note The parent physical interface passes untagged packets. You may not want to pass untagged packets, so be sure not to include the parent interface in your security policy.

Procedure

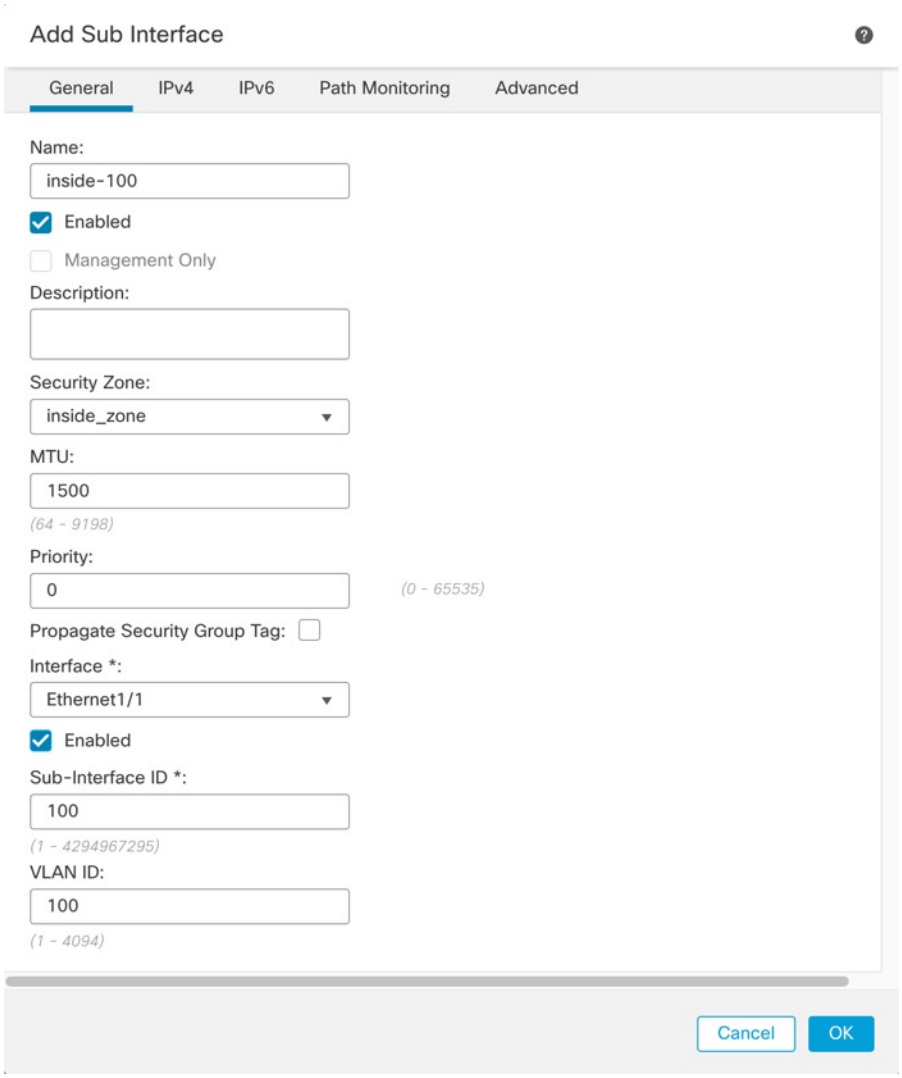
Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** () for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the parent interface according to Enable the physical interface and configure Ethernet settings, on page 1003.

Step 3 Click **Add Interfaces** > **Sub Interface**.

Step 4 On **General**, set these parameters:

Figure 256: Add subinterface



a) **Interface**—Choose the physical, redundant, or port-channel interface to which you want to add the subinterface.

b) **Sub-Interface ID**—Enter the subinterface ID as an integer between 1 and 4294967295. The number of subinterfaces allowed depends on your platform. You cannot change the ID after you set it.

c) **VLAN ID**—Enter the VLAN ID between 1 and 4094 that will be used to tag the packets on this subinterface.

   This VLAN ID must be unique.

Step 5 Click **OK**.

Step 6 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 7 Configure the routed or transparent mode interface parameters. See Configure routed mode interfaces, on page 1080 or Bridge group interfaces, on page 1086.

Configure VXLAN Interfaces

This chapter tells how to configure Virtual eXtensible LAN (VXLAN) interfaces. VXLAN interfaces act as Layer 2 virtual networks over Layer 3 physical networks to stretch Layer 2 networks.

VXLAN interfaces

Virtual Extensible Local Area Network (VXLAN) is a network encapsulation protocol that provides flexible placement of multitenant segments throughout the data center. VXLAN offers higher scalability, supporting up to 16 million Layer 2 segments. VXLAN delivers Ethernet Layer 2 network services similar to VLAN and provides greater extensibility and flexibility.

Additional information

For detailed information about VXLAN, see RFC 7348.

encapsulation

The Firewall Threat Defense supports two types of VXLAN encapsulation:

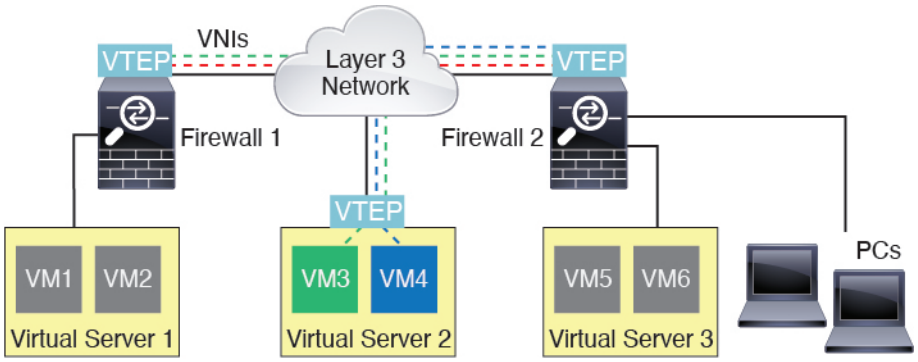
- VXLAN (all models)—VXLAN uses MAC Address-in-User Datagram Protocol (MAC-in-UDP) encapsulation. The original Layer 2 frame has a VXLAN header added and is then placed in a UDP-IP packet.
- Geneve (Firewall Threat Defense Virtual only)—Geneve has a flexible inner header that is not limited to the MAC address. Geneve encapsulation is required for transparent routing of packets between an Amazon Web Services (AWS) Gateway Load Balancer and appliances, and for sending extra information.

VXLAN tunnel endpoints

A VXLAN tunnel endpoint (VTEP) devices performs VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces to which you apply your security policy, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

VTEP network communication

This figure shows two Firewall Threat Defenses and Virtual Server 2 acting as VTEPs across a Layer 3 network, extending the VNI 1, 2, and 3 networks between sites. The Firewall Threat Defenses act as bridges or gateways between VXLAN and non-VXLAN networks.



The underlying IP network between VTEPs is independent of the VXLAN overlay. Encapsulated packets are routed based on the outer IP address header, which has the initiating VTEP as the source IP address and the terminating VTEP as the destination IP address. For VXLAN encapsulation: The destination IP address can be a multicast group when the remote VTEP is not known. With Geneve, the Firewall Threat Defense only supports static peers. The destination port for VXLAN is UDP port 4789 by default (user configurable). The destination port for Geneve is 6081.

VTEP source interface

A VTEP source interface is a regular interface (physical, EtherChannel, or even VLAN) that you plan to associate with all VNI interfaces. You can configure one VTEP source interface per Firewall Threat Defense Virtual.

VTEP source interface characteristics

The VTEP source interface can be devoted wholly to VXLAN traffic, although it is not restricted to that use. You can use the interface to handle regular traffic and apply a security policy to it. However, for VXLAN traffic, you must apply all security policies to the VNI interfaces. The VTEP interface serves as a physical port only.

In transparent firewall mode, the VTEP source interface is not part of a BVI. You must configure an IP address for it, similar to how the management interface is configured.

Because you can only configure one VTEP source interface, you cannot configure both VXLAN and Geneve interfaces on the same device. There is an exception for Firewall Threat Defense Virtual clustering on AWS or Azure, where you can have two VTEP source interfaces: a VXLAN interface is used for the cluster control link, and a Geneve (AWS) or VXLAN (Azure) interface can be used for the Gateway Load Balancer.

VNI interfaces

A VNI interface is a virtual interface that keeps network traffic separated on a given physical interface by using tagging. VNI interfaces are similar to VLAN interfaces. You can apply your security policy directly to each interface.

VNI interface configuration details

You can add only one VTEP interface. All VNI interfaces are associated with the same VTEP interface. There is an exception for Firewall Threat Defense Virtual clustering on AWS or Azure. For AWS clustering, you can have two VTEP source interfaces: a VXLAN interface is used for the cluster control link, and a Geneve interface can be used for the AWS Gateway Load Balancer. For Azure clustering, you can have two VTEP

source interfaces: a VXLAN interface is used for the cluster control link, and a second VXLAN interface can be used for the Azure Gateway Load Balancer.

How VXLAN and Geneve packets are processed

Summary

The key components involved in VXLAN and Geneve packet processing are:

- VTEP source interface: Performs encapsulation and decapsulation of VXLAN and Geneve traffic
- Inner MAC frame: The original packet payload that gets encapsulated
- Outer frame: The encapsulated packet with VXLAN or Geneve headers
- Remote VTEP: The destination endpoint for encapsulated traffic

Workflow

These stages describe how VXLAN and Geneve packet processing work:

1. Traffic entering and exiting the VTEP source interface is subject to VXLAN processing, specifically encapsulation or decapsulation.

   Encapsulation processing includes these tasks:

   - The VTEP source interface encapsulates the inner MAC frame with the VXLAN header.
   - The UDP checksum field is set to zero.
   - The Outer frame source IP is set to the VTEP interface IP.
   - The Outer frame destination IP is decided by a remote VTEP IP lookup.

   Decapsulation: the Firewall Threat Defense only decapsulates a VXLAN packet if:

   - It is a UDP packet with the destination port set to 4789 (this value is user configurable).
   - The ingress interface is the VTEP source interface.
   - The ingress interface IP address is the same as the destination IP address.
   - The VXLAN packet format is compliant with the standard.

2. Traffic entering and exiting the VTEP source interface is subject to Geneve processing, specifically encapsulation or decapsulation.

   Encapsulation processing includes these tasks:

   - The VTEP source interface encapsulates the inner MAC frame with the Geneve header.
   - The UDP checksum field is set to zero.
   - The Outer frame source IP is set to the VTEP interface IP.
   - The Outer frame destination IP is set the peer IP address that you configured.

   Decapsulation; the Firewall Threat Defense only decapsulates a Geneve packet if:

- It is a UDP packet with the destination port set to 6081 (this value is user configurable).
- The ingress interface is the VTEP source interface.
- The ingress interface IP address is the same as the destination IP address.
- The Geneve packet format is compliant with the standard.

Peer VTEPs

A peer VTEP is a VXLAN Tunnel Endpoint that enables the Firewall Threat Defense to send packets to devices behind remote VTEPs.

When the Firewall Threat Defense sends a packet to a device behind a peer VTEP, the Firewall Threat Defense needs two important pieces of information:

- The destination MAC address of the remote device
- The destination IP address of the peer VTEP

The Firewall Threat Defense maintains a mapping of destination MAC addresses to remote VTEP IP addresses for the VNI interfaces.

VXLAN peer configuration methods

There are two ways in which the Firewall Threat Defense can find this information:

- A single peer VTEP IP address can be statically configured on the Firewall Threat Defense.

  For IPv4: The Firewall Threat Defense then sends a VXLAN-encapsulated ARP broadcast to the VTEP to learn the end node MAC address.

  For IPv6: The Firewall Threat Defense then sends an IPv6 Neighbor Solicitation message to the IPv6 solicited-node multicast address. The peer VTEP responds with an IPv6 Neighbor Advertisement message with its link-local address.

- A group of peer VTEP IP addresses can be statically configured on the Firewall Threat Defense.

  For IPv4: The Firewall Threat Defense then sends a VXLAN-encapsulated ARP broadcast to the VTEP to learn the end node MAC addresses.

  For IPv6: The Firewall Threat Defense then sends an IPv6 Neighbor Solicitation message to the IPv6 solicited-node multicast address. The peer VTEP responds with an IPv6 Neighbor Advertisement message with its link-local address.

- A multicast group can be configured on each VNI interface (or on the VTEP as a whole).

  For IPv4: The Firewall Threat Defense sends a VXLAN-encapsulated ARP broadcast packet within an IP multicast packet through the VTEP source interface. The response to this ARP request enables the Firewall Threat Defense to learn both the remote VTEP IP address along with the destination MAC address of the remote end node.

  For IPv6: The Firewall Threat Defense sends a Multicast Listener Discovery (MLD) Report message through the VTEP source interface to indicate that the Firewall Threat Defense is listening on the VTEP interface for the multicast address traffic.

  This option is not supported with Geneve.

Geneve peer configuration

The Firewall Threat Defense Virtual only supports statically defined peers. You can define the Firewall Threat Defense Virtual peer IP address on the AWS Gateway Load Balancer. Because the Firewall Threat Defense Virtual never initiates traffic to the Gateway Load Balancer, you do not also have to specify the Gateway Load Balancer IP address on the Firewall Threat Defense Virtual. It learns the peer IP address when it receives Geneve traffic. Multicast groups are not supported with Geneve.

VXLAN use cases

This section describes the use cases for implementing VXLAN on the Firewall Threat Defense.

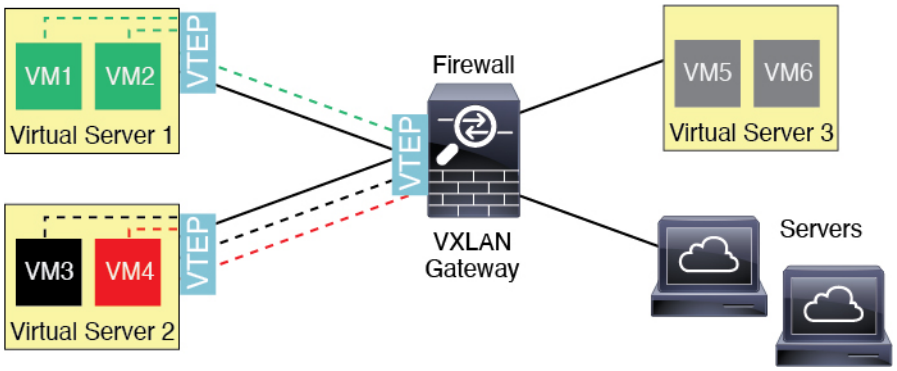
VXLAN bridge or gateway overview

A VTEP acts as a bridge or gateway between end nodes, such as VMs, servers, and PCs, and the VXLAN overlay network.

VXLAN Packet Processing

For incoming frames received with VXLAN encapsulation over the VTEP source interface, the Firewall Threat Defense strips out the VXLAN header and forwards it to a physical interface connected to a non-VXLAN network based on the destination MAC address of the inner Ethernet frame.

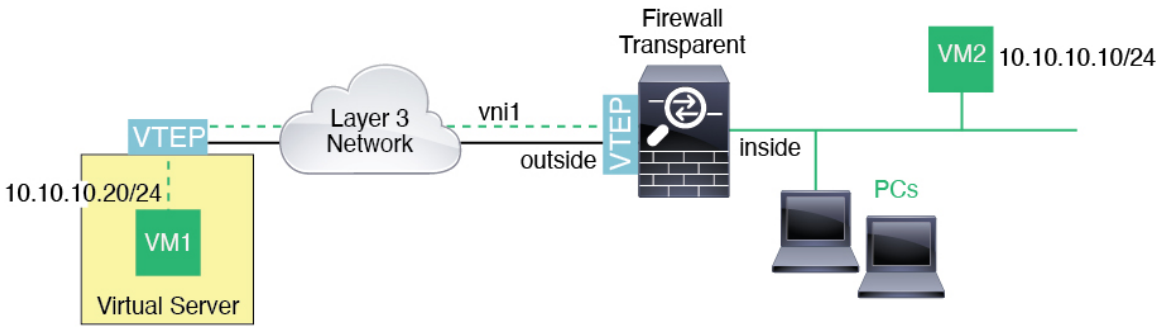
The Firewall Threat Defense always processes VXLAN packets rather than forwarding them untouched between two other VTEPs.



VXLAN bridge

A VXLAN bridge is a network configuration where the Firewall Threat Defense connects a remote VXLAN segment to a local segment.

When you use a bridge group (transparent firewall mode, or optionally routed mode), the Firewall Threat Defense can serve as a VXLAN bridge between a (remote) VXLAN segment and a local segment where both are in the same network. In this case, one member of the bridge group is a regular interface while the other member is a VNI interface.

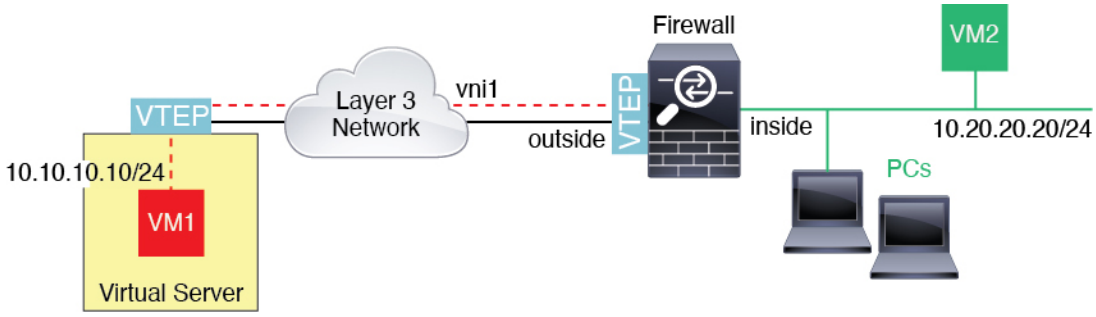


VXLAN gateway (routed mode)

The Firewall Threat Defense functions as a router between VXLAN and non-VXLAN domains, connecting devices on different networks.

VXLAN Gateway Visual Representation

This figure illustrates the VXLAN gateway connectivity.



Router Between VXLAN Domains

A VXLAN-stretched Layer 2 domain enables a VM to utilize a Firewall Threat Defense as its default gateway regardless of physical proximity.

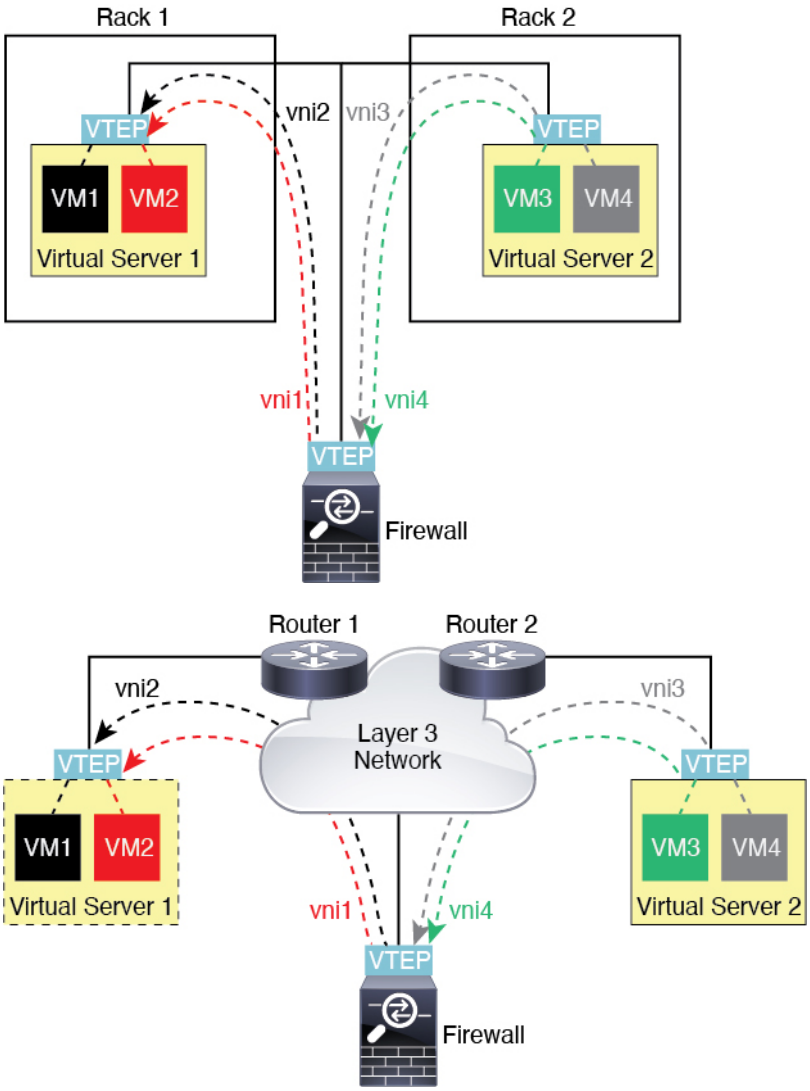
Routing Process and Packet Flow

The following steps outline the packet flow between VMs across VXLAN domains:

1. For packets from VM3 to VM1, the destination MAC address is the Firewall Threat Defense MAC address, because the Firewall Threat Defense is the default gateway.
2. The VTEP source interface on Virtual Server 2 receives packets from VM3, then encapsulates the packets with VNI 3's VXLAN tag and sends them to the Firewall Threat Defense .
3. When the Firewall Threat Defense receives the packets, it decapsulates the packets to get the inner frames.
4. The Firewall Threat Defense uses the inner frames for route lookup, then finds that the destination is on VNI 2. If it does not already have a mapping for VM1, the Firewall Threat Defense sends an encapsulated ARP broadcast on the multicast group IP on VNI 2.
5. The Firewall Threat Defense encapsulates the packets again with the VXLAN tag for VNI 2 and sends the packets to Virtual Server 1. Before encapsulation, the Firewall Threat Defense changes the inner frame destination MAC address to be the MAC of VM1.

6. When Virtual Server 1 receives the VXLAN packets, it decapsulates the packets and delivers the inner frames to VM1.

Note The Firewall Threat Defense must use dynamic VTEP peer discovery because it has multiple VTEP peers in this scenario.



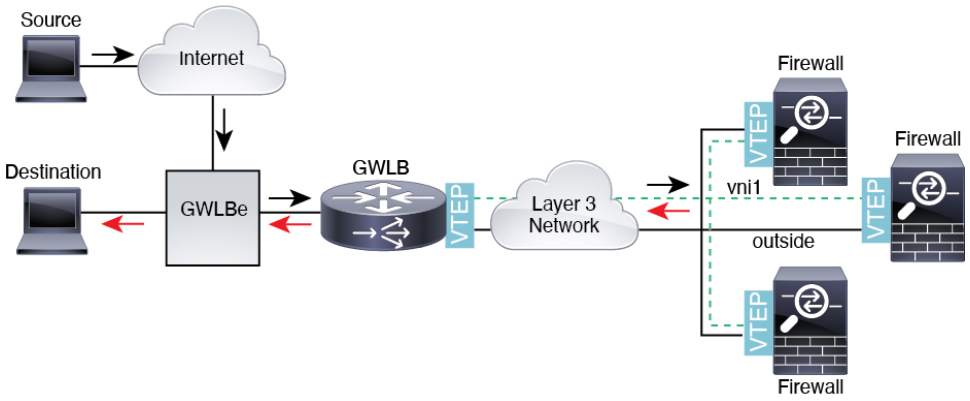
Geneve Single-Arm Proxy

In a Geneve single-arm proxy deployment, network traffic is routed through a centralized proxy device for inspection and policy enforcement.

Note This use case is the only currently supported use case for Geneve interfaces.

The AWS Gateway Load Balancer combines a transparent network gateway and a load balancer that distributes traffic and scales virtual appliances on demand. The Firewall Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint). The following figure shows traffic forwarded to the Gateway Load Balancer from the Gateway Load Balancer endpoint. The Gateway Load Balancer balances traffic among multiple Firewall Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer (U-turn traffic). The Gateway Load Balancer then sends the traffic back to the Gateway Load Balancer endpoint and to the destination.

Figure 257: Geneve Single-Arm Proxy

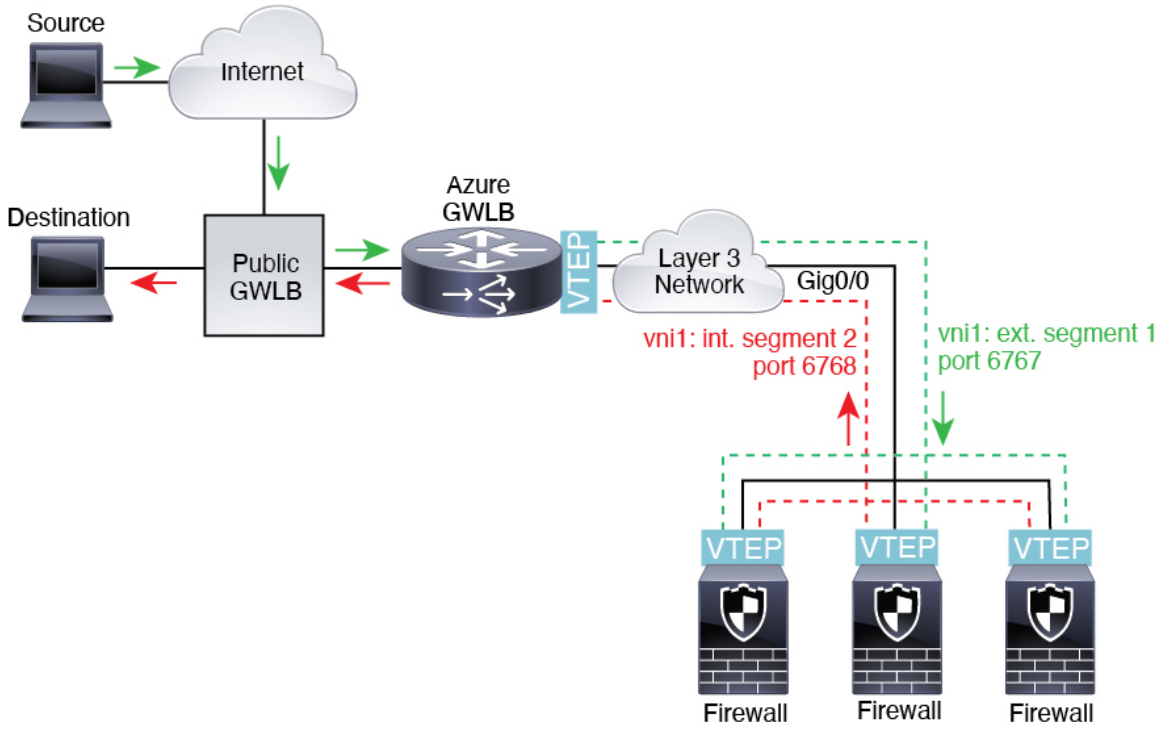


Azure Gateway Load Balancer and Paired Proxy

In an Azure service chain, Firewall Threat Defense Virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The Firewall Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy.

The following figure shows traffic forwarded to the Azure Gateway Load Balancer from the Public Gateway Load Balancer on the external VXLAN segment. The Gateway Load Balancer balances traffic among multiple Firewall Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer on the internal VXLAN segment. The Azure Gateway Load Balancer then sends the traffic back to the Public Gateway Load Balancer and to the destination.

Figure 258: Azure Gateway Load Balancer with Paired Proxy



Requirements and prerequisites for VXLAN interfaces

This reference provides the model requirements and prerequisites for configuring VXLAN interfaces on supported platforms.

Model requirements

- VXLAN encapsulation is supported on all models.
- Geneve encapsulation is supported for these models:
  - Firewall Threat Defense Virtual in Amazon Web Services (AWS)
- VXLAN in *paired proxy mode* is supported for these models:
  - Firewall Threat Defense Virtual in Azure
- Firepower 1010 and Secure Firewall 1210/1220 switch ports and VLAN interfaces are not supported as VTEP interfaces.

Guidelines for VXLAN interfaces

Firewall mode restrictions

Geneve interfaces are only supported in routed firewall mode.

- Paired proxy VXLAN interfaces are only supported in routed firewall mode.

IPv6 support restrictions

The VNI interface supports both IPv4 and IPv6 traffic.

- For VXLAN encapsulation, the VTEP source interface supports both IPv4 and IPv6. The Firewall Threat Defense Virtual cluster control link VTEP source interface only supports IPv4.

  For Geneve, the VTEP source interface only supports IPv4.

Clustering restrictions

Clustering does not support VXLAN in Individual Interface mode except for the cluster control link (Firewall Threat Defense Virtual only) Only Spanned EtherChannel mode supports VXLAN.

An exception is made for AWS, which can use an additional Geneve interface for use with the GWLB and for Azure, which can use an additional paired proxy VXLAN interface for use with the GWLB.

Routing restrictions

Only static routing or Policy Based Routing is supported on the VNI interface; dynamic routing protocols are not supported.

VPN restrictions

You cannot configure the VTEP source interface for VPN or use it as a VTI.

MTU requirements

If the source interface MTU is less than the required minimum, the Firewall Threat Defense automatically raises the MTU to meet encapsulation requirements.

- VXLAN encapsulation—If the source interface MTU is less than 1554 bytes for IPv4 or 1574 bytes for IPv6, then the Firewall Threat Defense automatically raises the MTU to 1554 bytes or 1574 bytes. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. If the MTU used by other devices is larger, then you should set the source interface MTU to be the network MTU + 54 bytes for IPv4 or +64 bytes for IPv6. For the Firewall Threat Defense Virtual, this MTU requires a restart to enable jumbo frame reservation.
- Geneve encapsulation—If the source interface MTU is less than 1806 bytes, then the Firewall Threat Defense automatically raises the MTU to 1806 bytes. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. If the MTU used by other devices is larger, then you should set the source interface MTU to be the network MTU + 306 bytes. This MTU requires a restart to enable jumbo frame reservation.

Configure VXLAN or Geneve interfaces

You can configure either VXLAN or Geneve interfaces.

Configure VXLAN interfaces

Configure VXLAN interfaces to enable virtualized network overlay functionality and network segmentation.

VXLAN provides Layer 2 overlay network capabilities over existing Layer 3 infrastructure. You can configure either VXLAN or Geneve interfaces, but not both simultaneously.

Note You can configure either VXLAN or Geneve (Firewall Threat Defense Virtual only). For Geneve interfaces, see Configure geneve interfaces, on page 1074.

Note For Azure GWLB, the VXLAN interface is configured when you deploy the VM using the ARM template. You can use this section to change your configuration.

Follow these steps to configure VXLAN interfaces.

Procedure

Step 1 Configure the VTEP source interface, on page 1072.

Step 2 Configure the VNI interface, on page 1073.

Step 3 (Azure GWLB) Allow gateway load balancer health checks, on page 1076.

Configure the VTEP source interface

This task enables you to set up VXLAN network virtualization by configuring a Network Virtualization Endpoint (NVE) on your firewall device.

You can configure one VTEP source interface per Firewall Threat Defense device. The VTEP is defined as a Network Virtualization Endpoint (NVE). VXLAN is the default encapsulation type. An exception is made for clustering on the Firewall Threat Defense Virtual in Azure, where you can use one VTEP source interface for the cluster control link and a second one for the data interface connected to the Azure GWLB.

Procedure

Step 1 If you want to specify a group of peer VTEPs, add a network object with the peer IP addresses. See Create a network object, on page 2985.

Step 2 Choose **Devices** > **Device Management**.

Step 3 Click **Edit** (✎) next to the device on which you want to configure VXLAN.

Step 4 (Optional) Specify that the source interface is NVE-only.

This setting is optional for routed mode where this setting restricts traffic to VXLAN and common management traffic only on this interface. This setting is automatically enabled for transparent firewall mode.

a) Click **Interfaces**.

b) Click **Edit** (✎) for the VTEP source interface.

c) On the **General** page, check the check box of **NVE Only**.

Step 5 Click **VTEP** if it is not already displaying.

Step 6 Check **Enable NVE**.

Step 7 Click **Add VTEP**.

a) For the **Encapsulation Type**, choose **VXLAN**.

For AWS, you can choose between **VXLAN** and **Geneve**. Other platforms have **VXLAN** chosen automatically.

b) Enter the value for the **Encapsulation port** within the specified range.

The default value is 4789.

c) Select the **VTEP Source Interface**.

Select from the list of available physical interfaces present on the device. If the source interface MTU is less than 1554 bytes for IPv4 or 1574 bytes for IPv6, then the Cloud-Delivered Firewall Management Center automatically raises the MTU to 1554 bytes or 1574 bytes.

d) Select the **Neighbor Address**. The available options are:

- **None**—No neighbor address is specified.
- **Peer VTEP**—Specify a peer VTP address.
- **Peer Group**—Specify a network object with the peer IP addresses.
- **Default Multicast**—Specify a default multicast group for all associated VNI interfaces. If you do not configure the multicast group per VNI interface, then this group is used. If you configure a group at the VNI interface level, then that group overrides this setting.

Step 8 Click **OK**.

Step 9 Click **Save**.

What to do next

Configure the routed interface parameters. See Configure Routed Mode Interfaces.

Configure the VNI interface

Add a VNI interface, associate it with the VTEP source interface, and configure basic interface parameters.

For the Firewall Threat Defense Virtual in Azure, you can configure either a regular VXLAN interface, or you can configure a paired proxy mode VXLAN interface for use with the Azure GWLB. Paired proxy mode is the only supported mode with clustering.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device on which you want to configure VXLAN.

Step 3 Click **Interfaces**.

Step 4 Click **Add Interfaces**, and then choose **VNI Interface**.

a) Enter the interface **Name** and **Description**.
b) From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.
c) Enter a value for the **Priority** field within the specified range. By default, 0 is selected.
d) Enter a value for the **Priority** field within the specified range. By default, 0 is selected.
e) Enter a value for the **VNI ID** between 1 and 10000.

This ID is only an internal interface identifier.

Step 5 (Paired Proxy VXLAN for Azure gateway load balancer) Enable proxy paired mode and set the required parameters.

a) Check **Proxy Paired**.
b) Set the **Internal Port** between 1024 and 65535.
c) Set the **Internal Segment ID** between 1 and 16777215.
d) Set the **External Port** between 1024 and 65535.
e) Set the **External Segment ID** between 1 and 16777215.

Step 6 (Regular VXLAN) Enter a value for the **VNI Segment ID** between 1 and 16777215.

The segment ID is used for VXLAN tagging.

Step 7 Enter the **Multicast Group IP Address**.

If you do not set the multicast group for the VNI interface, the default group from the VTEP source interface configuration is used, if available. If you manually set a VTEP peer IP for the VTEP source interface, you cannot specify a multicast group for the VNI interface.

Step 8 Check **NVE Mapped to VTEP Interface**.

This option associates this interface with the VTEP source interface.

Step 9 Click **OK**.

Click **Save** to save the interface configuration.

What to do next

Configure the routed or transparent interface parameters. See Configure routed and transparent mode interfaces, on page 1077.

Configure geneve interfaces

To configure Geneve interfaces for Firewall Threat Defense Virtual, perform the following steps.

Note You can configure either VXLAN or Geneve. For information about VXLAN interfaces, see Configure VXLAN interfaces, on page 1071.

1. Configure the VTEP source interface, on page 1075.
2. Configure the VNI, on page 1075.
3. Allow gateway load balancer health checks, on page 1076.

Configure the VTEP source interface

Configure the VTEP source interface to enable network virtualization functionality on the device.

You can configure one VTEP source interface per Firewall Threat Defense Virtual device. The VTEP is defined as a Network Virtualization Endpoint (NVE).

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device on which you want to configure Geneve.

Step 3 Click **VTEP**.

Step 4 Check **Enable NVE**.

Step 5 Click **Add VTEP**.

a) For the **Encapsulation Type**, choose **Geneve**.

b) Enter the value for the **Encapsulation port** within the specified range.

We do not recommend changing the Geneve port; AWS requires a port of 6081.

c) Select the **VTEP Source Interface**.

You can select from the list of available physical interfaces present on the device. If the source interface MTU is less than 1806 bytes, then the Cloud-Delivered Firewall Management Center automatically raises the MTU to 1806 bytes.

d) Click **OK**.

Step 6 Click **Save**.

Step 7 Configure the routed interface parameters. See Configure Routed Mode Interfaces.

Configure the VNI

This task enables you to add a VNI interface and configure its basic parameters for Geneve functionality on your device.

Add a VNI, associate it with (VTEP) source interface, and configure basic interface parameters.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** ( ) next to the device on which you want to configure Geneve.

Step 3 Click **Interfaces**.

Step 4 Click **Add Interfaces**, and then choose **VNI Interface**.

a) In the **Name** and **Description** fields, provide relevant information.

b) In the **VNI ID** field, enter a value between 1 and 10000.

Note
This ID is only an internal interface identifier.

Step 5 Check the **Enable Proxy** check box.

Note
When the AWS proxy is enabled on the VNI interface of your device, configuring NAT is not allowed.

This option enables single-arm proxy, and allows traffic to exit the same interface it entered (U-turn traffic). If you edit the interface later , you cannot disable single-arm proxy. To disable single-arm proxy, delete the existing interface and create a new VNI interface.

This option is only available for a Geneve VTEP.

Step 6 Select **NVE Mapped to VTEP Interface**.

This option associates this interface with the VTEP source interface.

Step 7 Click **OK**.

Step 8 Click **Save**.

What to do next

Configure the routed interface parameters. See Configure Routed Mode Interfaces.

Allow gateway load balancer health checks

The AWS or Azure gateway load balancer (GWLB) requires appliances to answer a health check properly. The GWLB will only send traffic to appliances that are considered healthy.

Configure one of these methods to enable proper health check responses.

Procedure

Step 1 Configure SSH. See Configure Secure Shell

Allow SSH from the GWLB IP address. The GWLB will attempt to establish a connection to the Firewall Threat Defense Virtual, and the Firewall Threat Defense Virtual's prompt to log in is taken as proof of health. An SSH login attempt will time out after 1 minute. You will need to configure a longer health check interval on the GWLB to accommodate this timeout.

Step 2 Configure HTTP(S) Redirection Using Static Interface NAT with Port Translation.

You can configure the Firewall Threat Defense Virtual to redirect health checks to a metadata HTTP(S) server. For HTTP(S) health checks, the HTTP(S) server must reply to the GWLB with a status code in the range 200 to 399. Because the Firewall Threat Defense Virtual has limits on the number of simultaneous management connections, you may choose to offload the health check to an external server.

Static interface NAT with port translation lets you redirect a connection to a port (such as port 80) to a different IP address. For example, translate an HTTP packet from the GWLB with a destination of the Firewall Threat Defense Virtual outside interface so that it appears to be from the Firewall Threat Defense Virtual outside interface with a destination of the HTTP server. The Firewall Threat Defense Virtual then forwards the packet to the mapped destination address. The HTTP server responds to the Firewall Threat Defense Virtual outside interface, and then the Firewall Threat Defense Virtual forwards the response back to the GWLB. You need an access rule that allows traffic from the GWLB to the HTTP server.

a) Permit HTTP(S) traffic on the outside interface from the GWLB network in an access rule. See Access Control Rules, on page 2041.
b) For HTTP(S), translate the source GWLB IP address to the Firewall Threat Defense Virtual outside interface IP address; then translate the destination of the outside interface IP address to the HTTP(S) server IP address. See Configure static manual NAT, on page 1614.

Configure routed and transparent mode interfaces

This section includes tasks to complete the regular interface configuration for all models in routed or transparent firewall mode.

Routed and transparent mode interfaces

Firewall mode interfaces subject traffic to firewall functions such as maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization. You can also optionally configure IPS functions for this traffic according to your security policy.

Interface types by firewall mode

The types of firewall interfaces you can configure depends on the firewall mode set for the device: routed or transparent mode. See Transparent or Routed Firewall Mode, on page 153 for more information.

- Routed mode interfaces (routed firewall mode only)—Each interface that you want to route between is on a different subnet.
- Bridge group interfaces (routed and transparent firewall mode)—You can group together multiple interfaces on a network, and the Firewall Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. In routed mode, the Firewall Threat Defense device routes between BVIs and regular routed interfaces. In transparent mode, each bridge group is separate and cannot communicate with each other.

Dual IP stack (IPv4 and IPv6)

A dual IP stack is a network configuration that enables both IPv6 and IPv4 addresses on the same interface on Firewall Threat Defense device. Make sure you configure a default route for both IPv4 and IPv6.

31-bit subnet masks

A 31-bit subnet mask includes only two addresses for point-to-point connections.

31-bit subnet mask implementation

For routed interfaces, you can configure an IP address on a 31-bit subnet for point-to-point connections. Normally, the first and last address in a 2-address subnet is reserved for the network and broadcast, so a 2-address subnet is not usable. However, if you have a point-to-point connection and do not need network or broadcast addresses, a 31-bit subnet is a useful way to preserve addresses in IPv4.

31-bit subnet mask use cases

The failover link between two Firewall Threat Defenses only requires two addresses. Any packet that is transmitted by one end of the link is always received by the other, and broadcasting is unnecessary. You can also have a directly-connected management station running SNMP or Syslog.

Note: 31-bit subnet and clustering compatibility

You can use a 31-bit subnet mask for cluster interfaces, excluding the management interface and the cluster control link.

31-bit subnet and failover

A Firewall Threat Defense using a 31-bit subnet for its interface IP address can't configure a standby IP address due to insufficient addresses.

Interface monitoring limitations

When using a 31-bit subnet for the Firewall Threat Defense interface IP address, you cannot configure a standby IP address for the interface because there are not enough addresses. Normally, an interface for failover should have a standby IP address so the active unit can perform interface tests to ensure standby interface health. Without a standby IP address, the Firewall Threat Defense cannot perform any network tests; only the link state can be tracked.

For the failover and optional separate state link, which are point-to-point connections, you can also use a 31-bit subnet.

31-bit subnet and management

If you have a directly-connected management station, you can use a point-to-point connection for SSH or HTTP on the Firewall Threat Defense, or for SNMP or Syslog on the management station.

Recommendation: 31-bit subnet unsupported features

The 31-bit subnet does not support certain network features that require multiple host addresses.

- BVI interfaces for bridge groups—The bridge group requires at least 3 host addresses: the BVI, and two hosts connected to two bridge group member interfaces. you must use a /29 subnet or smaller.
- Multicast Routing

Guidelines and limitations for routed and transparent mode interfaces

Context mode

- PPPoE is not supported in multiple context mode.
- For multiple context mode in transparent mode, each context must use different interfaces; you cannot share an interface across contexts.
- For multiple context mode in transparent mode, each context typically uses a different subnet. You can use overlapping subnets, but your network topology requires router and NAT configuration to make it possible from a routing standpoint.

High availability, Clustering, and Multi-Instance

Do not configure failover links with the procedures in this chapter. See the High availability chapter for more information.

- For cluster interfaces, see the clustering chapter for requirements.
- For multi-instance mode, shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode).
- When you use High availability, you must set the IP address and standby address for data interfaces manually; DHCP and PPPoE are not supported. Set the standby IP addresses on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. See the High availability chapter for more information.

IPv6

IPv6 is supported on all interfaces.

- You can only configure IPv6 addresses manually in transparent mode.
- The Firewall Threat Defense device does not support IPv6 anycast addresses.
- DHCPv6 and prefix delegation options are not supported with transparent mode, clustering, or High availability.

Model guidelines

For the Firewall Threat Defense Virtual on VMware with bridged ixgbevf interfaces, bridge groups are not supported.

- For the Firepower 2100 series, bridge groups are not supported in routed mode.

VLAN IDs for the ASASM

- If you add an interface for a VLAN that is not yet assigned to the ASA by the switch, the interface will be in the down state. When you assign the VLAN to the ASA, the interface changes to an up state. See the **show interface** command for more information about interface states.

Transparent mode and bridge group guidelines

- You can create up to 250 bridge groups, with 64 interfaces per bridge group.
- Each directly-connected network must be on the same subnet.
- The Firewall Threat Defense device does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.
- An IP address for the BVI is required for each bridge group for to-the-device and from-the-device management traffic, as well as for data traffic to pass through the Firewall Threat Defense device. For IPv4 traffic, specify an IPv4 address. For IPv6 traffic, specify an IPv6 address.
- You can only configure IPv6 addresses manually.
- The BVI IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255).

- Management interfaces are not supported as bridge group members.
- For multi-instance mode, shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode).
- For the Firewall Threat Defense Virtual on VMware with bridged ixgbevf interfaces, transparent mode is not supported, and bridge groups are not supported in routed mode.
-
- For the 200/ 1010/ 1210// 1220, you cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.
- For the Firepower 4100/9300, data-sharing interfaces are not supported as bridge group members.
- In transparent mode, you must use at least 1 bridge group; data interfaces must belong to a bridge group.
- In transparent mode, do not specify the BVI IP address as the default gateway for connected devices; devices need to specify the router on the other side of the Firewall Threat Defense as the default gateway.
- In transparent mode, the *default* route, which is required to provide a return path for management traffic, is only applied to management traffic from one bridge group network. This is because the default route specifies an interface in the bridge group as well as the router IP address on the bridge group network, and you can only define one default route. If you have management traffic from more than one bridge group network, you need to specify a regular static route that identifies the network from which you expect management traffic.
- Transparent mode is not supported on threat defense virtual instances deployed on Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud Infrastructure.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, Firewall Threat Defense-defined EtherChannel interfaces are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.
- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the Firewall Threat Defense when using bridge group members. If there are two neighbors on either side of the Firewall Threat Defense running BFD, then the Firewall Threat Defense will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

Additional guidelines and requirements

The Firewall Threat Defense supports only one 802.1Q header in a packet and does not support multiple headers (known as Q-in-Q support) for firewall interfaces.**Note:** For inline sets and passive interfaces, the FTD supports Q-in-Q up to two 802.1Q headers in a packet, with the exception of the Firepower 4100/9300, which only supports one 802.1Q header.

- Interface problems, such as frequent up/down status changes, can prevent the floating connection timer from applying correctly to the connections going through the interface. If you have problems with an interface's status, consider clearing all connections after the status becomes stable to clear invalid connections.

Configure routed mode interfaces

Before you begin

- Firepower 4100/9300
  1. Configure a physical interface
  2. (Optional) Configure any special interfaces.
     - Add an EtherChannel (port channel)
     - Add a VLAN subinterface for container instances in FXOS
     - Configure a loopback interface, on page 1053
     - Add a subinterface, on page 1060 in Cloud-Delivered Firewall Management Center
     - Configure VXLAN interfaces, on page 1071

- (Optional) **All other models:**
  - Configure an EtherChannel, on page 1011
  - Configure a loopback interface, on page 1053
  - Add a subinterface, on page 1060
  - Configure VXLAN interfaces, on page 1071
  - Firewall Threat Defense Virtual on AWS: Configure geneve interfaces, on page 1074
  - Firepower 1010 and Secure Firewall 1210/1220 : Configure a VLAN interface, on page 1044

This procedure describes how to set the name, security zone, and IPv4 address.

Note Not all fields are supported for all interface types.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 In the **Name** field, enter a name up to 48 characters in length.

You cannot start the name with the phrase "cluster". It is reserved for internal use.

Step 4 Enable the interface by checking the **Enabled** check box.

a) (Optional) Set this interface to **Management Only** to limit traffic to management traffic; through-the-box traffic is not allowed.

b) (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

c) In the **Mode** drop-down list, choose **None** .

Regular firewall interfaces have the mode set to None. The other modes are for IPS-only interface types.

d) From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New** .

The routed interface is a Routed-type interface, and can only belong to Routed-type zones.

e) See Configure the MTU, on page 1109 for information about the **MTU** .

f) In the **Priority** field, enter a number ranging from 0–65535.

This value is used in the policy based routing configuration. The priority is used to determine how you want to route the traffic across multiple egress interfaces. For more information, see Configure policy-based routing policy, on page 1548.

Step 5 Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

High Availability, clustering , and loopback interfaces only support static IP address configuration; DHCP and PPPoE are not supported.

- **Use Static IP** —Enter the IP address and subnet mask. For point-to-point connections, you can specify a 31-bit subnet mask (255.255.255.254 or /31). In this case, no IP addresses are reserved for the network or broadcast addresses. You cannot set the standby IP address in this case. For High Availability, you can only use a static IP address. Set the standby IP address on the **Devices** > **Device Management** page, click **High Availability** in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.
- **Use DHCP** —Configure the following optional parameters:
  - **Obtain default route using DHCP** —Obtains the default route from the DHCP server.
  - **DHCP route metric** —Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.
- **Use PPPoE** —If the interface is connected to a DSL, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address, configure the following parameters:
  - **VPDN Group Name** —Specify a group name of your choice to represent this connection.
  - **PPPoE User Name** —Specify the username provided by your ISP.
  - **PPPoE Password/Confirm Password** —Specify and confirm the password provided by your ISP.
  - **PPP Authentication** —Choose **PAP** , **CHAP** , or **MSCHAP** .

    PAP passes a cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.
  - **PPPoE route metric** —Assign an administrative distance to the learned route. Valid values are from 1 to 255. By default, the administrative distance for the learned routes is 1.
  - **Enable Route Settings** —To manually configure the PPPoE IP address, check this box and then enter the **IP Address** .

    If you select the **Enable Route Settings** check box and leave the **IP Address** blank, the **ip address pppoe setroute** command is applied as shown in this example:

```
interface GigabitEthernet0/2
nameif inside2_pppoe
cts manual
  propagate sgt preserve-untag
  policy static sgt disabled trusted
security-level 0
pppoe client vpdn group test
pppoe client route distance 10
ip address pppoe setroute
```

- **Store Username and Password in Flash** —Stores the username and password in flash memory.

  The Firewall Threat Defense device stores the username and password in a special location of NVRAM.

Step 6 (Optional) Configure these fields, if you require them.

a) (Optional) See Configure IPv6 addressing, on page 1090 to configure IPv6 addressing on the **IPv6** tab.

b) (Optional) See Configure the MAC address, on page 1110 to manually configure the MAC address on the **Advanced** tab.

c) Enable Cloud-Delivered Firewall Management Center manager access on a data interface on the **Manager Access** page.

You can enable manager access from a data interface when you first setup the Firewall Threat Defense . If you want to enable or disable manager access after you added the Firewall Threat Defense to the Cloud-Delivered Firewall Management Center , see:

- Enable manager access: Change the manager access interface from management to data, on page 1174

  Note
  You cannot enable manager access unless you first initiate the manager interface migration from Management to a data interface. After you initiate the migration, you can enable manager access on the **Manager Access** page and save the configuration successfully.

- Disable manager access: Change the manager access interface from data to management, on page 1178

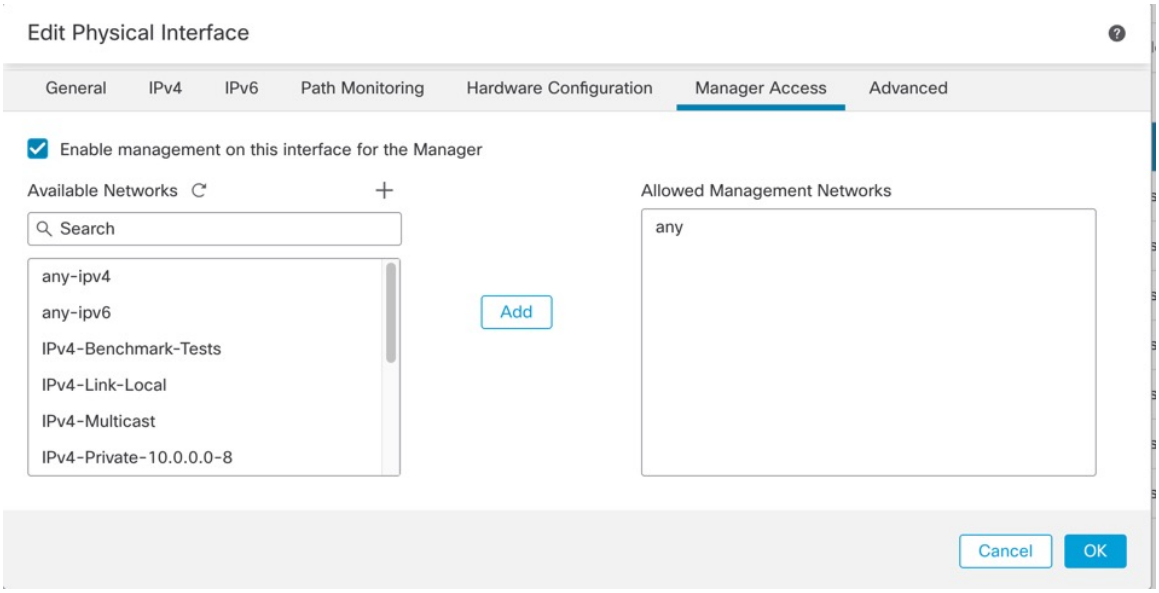
If you want to change the manager access interface from one data interface to another data interface, you must disable manager access on the original data interface, but do not disable the interface itself yet; the original data interface must be used to perform the deployment. If you want to use the same IP address on the new manager access interface, you can delete or change the IP configuration on the original interface; this change should not affect the deployment. If you use a different IP address for the new interface, then also change the device IP address shown in the Cloud-Delivered Firewall Management Center ; see Update the Hostname or IP Address in the Cloud-Delivered Firewall Management Center, on page 1167 . Be sure to also update related configuration to use the new interface such as static routes, DDNS, and DNS settings.

Manager access from a data interface has the following limitations:

- You can only enable manager access on a physical, data interface. You cannot use a subinterface or EtherChannel, nor can you create a subinterface on the manager access interface. You can also use the Cloud-Delivered Firewall Management Center to enable manager access on a single secondary interface for redundancy.
- This interface cannot be management-only.
- Routed firewall mode only, using a routed interface.

- PPPoE is not supported. If your ISP requires PPPoE, you will have to put a router with PPPoE support between the Firewall Threat Defense and the WAN modem.
- The interface must be in the global VRF only.
- SSH is not enabled by default for data interfaces, so you will have to enable SSH later using the Cloud-Delivered Firewall Management Center. Because the Management interface gateway will be changed to be the data interfaces, you also cannot SSH to the Management interface from a remote network unless you add a static route for the Management interface using the **configure network static-routes** command. For Firewall Threat Defense Virtual on Amazon Web Services, a console port is not available, so you should maintain your SSH access to the Management interface: add a static route for Management before you continue with your configuration. Alternatively, be sure to finish all CLI configuration (including the **configure manager add** command) before you configure the data interface for manager access and you are disconnected.
- Clustering is not supported. You must use the Management interface in this case.

Figure 259: Manager Access



- Check **Enable management on this interface for the manager** to use this data interface for management instead of the dedicated Management interface.
- (Optional) In the **Allowed Management Networks** box, add the networks from which you want to allow manager access. By default, any networks are allowed.

Step 7 (Optional) Set the duplex and speed by clicking **Hardware Configuration** > **Speed**.

- **Duplex**—Choose **Full** or **Half**. SFP interfaces only support **Full** duplex. However, the SFP interface (Ethernet 1/5) of Secure Firewall 220 supports **Full** and **Half** duplex.
- **Speed**—Choose a speed (varies depending on the model).For SFPs, choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always full, autonegotiation is enabled, and FEC is set to auto. For dual-speed transceivers, the lower speed is used. This option is useful if you later change the network module to a different model, and want the speed to update

automatically. Some switches and transceivers don't support autonegotiation, especially for higher speed interfaces. In this case, set the interface on the Firewall Threat Defense to a manual speed and also disable **Auto-negotiation** so the link can come up.

Note
You cannot modify the speed of a high availability or a cluster control link interface.

- **Auto-negotiation**—Set the interface to negotiate the link status and flow control.

  Except for the 1100 and 2100, autonegotiation is set separately from the speed. Some switches don't support autonegotiation, especially for higher speed interfaces. In this case, set the interface on the Firewall Threat Defense to a manual speed and also disable **Auto-negotiation** so the link can come up.

- **Forward Error Correction Mode**—For 25Gbps and higher interfaces, enable Forward Error Correction (FEC).

  For an EtherChannel member interface, you must configure FEC before you add it to the EtherChannel. If you remove the interface from the EtherChannel, after rebooting, you may need to reconfigure the FEC for the interface.

  Some switches don't support auto mode for FEC, especially for larger interfaces. Be sure to manually configure the setting, depending on the switch support. We don't recommend disabling FEC; if you do so, the EtherChannel may not come up properly.

  The setting chosen when you use **auto** depends on the transceiver type and whether the interface is fixed (built-in) or on a network module.

Table 67: Default FEC for Auto Setting

Transceiver Type	Fixed Port Default FEC (Ethernet 1/9 through 1/16)	Network Module Default FEC
25G-SR	Clause 74 FC-FEC	Clause 108 RS-FEC
25G-LR	Clause 74 FC-FEC	Clause 108 RS-FEC
10/25G-CSR	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-AOC*x*M	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-CU2.5/3M	Auto-Negotiate	Auto-Negotiate
25G-CU4/5M	Auto-Negotiate	Auto-Negotiate
25/50/100G	Clause 91 RS-FEC	Clause 91 RS-FEC

Note
From Version 7.4.2, Firewall Threat Defense supports the SFP 10G_25G_CSR_S module. To avoid link-down issues after a reboot, Forward Error Correction (FEC) must be enabled on these interfaces.

Step 8 Click **OK** .

Step 9 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Bridge group interfaces

A bridge group is a group of interfaces that the Secure Firewall Threat Defense device bridges instead of routes.

Bridge group support

Bridge groups are supported in both transparent and routed firewall mode.

For more information about bridge groups, see Bridge groups, on page 155.

To configure bridge groups and associated interfaces, perform these steps.

Configure general bridge group member interface parameters

This task configures the name and security zone for each bridge group member interface to enable bridge group functionality.

The same bridge group can include different types of interfaces: physical interfaces, VLAN subinterfaces, Firepower 1010 and Secure Firewall 1210/1220 VLAN interfaces, EtherChannels, and redundant interfaces. The Management interface is not supported. In routed mode, EtherChannels are not supported. For the Firepower 4100/9300, data-sharing type interfaces are not supported.

Before you begin

- Firepower 4100/9300
  1. Configure a physical interface
  2. (Optional) Configure any special interfaces.
     - Add an EtherChannel (port channel)
     - Add a VLAN subinterface for container instances in FXOS
     - Add a subinterface, on page 1060 in Cloud-Delivered Firewall Management Center
- (Optional) **All other models:**
  - Configure an EtherChannel, on page 1011
  - Add a subinterface, on page 1060
  - Firepower 1010 and Secure Firewall 1210/1220: Configure a VLAN interface, on page 1044

Follow these steps to configure general bridge group member interface parameters:

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Configure these fields:

a) In the **Name** field, enter a name up to 48 characters in length.

You cannot start the name with the phrase "cluster". It is reserved for internal use.

b) Enable the interface by checking the **Enabled** check box.

c) (Optional) Set this interface to **Management Only** to limit traffic to management traffic; through-the-box traffic is not allowed.

d) (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

e) In the **Mode** drop-down list, choose **None**.

Regular firewall interfaces have the mode set to None. The other modes are for IPS-only interface types. After you assign this interface to a bridge group, the mode will show as **Switched**.

f) From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

The bridge group member interface is a Switched-type interface, and can only belong to Switched-type zones. Do not configure any IP address settings for this interface. You will set the IP address for the Bridge Virtual Interface (BVI) only. Note that the BVI does not belong to a zone, and you cannot apply access control policies to the BVI.

Step 4 Refer to Configure the MTU, on page 1109 for information about **MTU**.

Step 5 (Optional) Set the duplex and speed by clicking **Hardware Configuration** > **Speed** .

- **Duplex** —Choose **Full** or **Half** . SFP interfaces only support **Full** duplex. However, the SFP interface (Ethernet 1/5) of Secure Firewall 220 supports **Full** and **Half** duplex.
- **Speed** —Choose a speed (varies depending on the model). For SFPs, choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always full, autonegotiation is enabled, and FEC is set to auto. For dual-speed transceivers, the lower speed is used. This option is useful if you later change the network module to a different model, and want the speed to update automatically. Some switches and transceivers don't support autonegotiation, especially for higher speed interfaces. In this case, set the interface on the Firewall Threat Defense to a manual speed and also disable **Auto-negotiation** so the link can come up.

  Note
  You cannot modify the speed of a high availability or a cluster control link interface.

- **Auto-negotiation** —Set the interface to negotiate the link status and flow control.

  Except for the 1100 and 2100 , autonegotiation is set separately from the speed. Some switches don't support autonegotiation, especially for higher speed interfaces. In this case, set the interface on the Firewall Threat Defense to a manual speed and also disable **Auto-negotiation** so the link can come up.

- **Forward Error Correction Mode** —For 25Gbps and higher interfaces, enable Forward Error Correction (FEC).

  For an EtherChannel member interface, configure FEC before adding it to the EtherChannel. If you remove the interface from EtherChannel and then reboot, reconfigure the FEC for the interface.

  Some switches don't support auto mode for FEC, especially for larger interfaces. Be sure to manually configure the setting, depending on the switch support. We don't recommend disabling FEC because the EtherChannel may not function correctly.

  The setting chosen when you use **auto** depends on the transceiver type and whether the interface is fixed (built-in) or on a network module.

  Table 68: Default FEC for Auto Setting

Transceiver Type	Fixed Port Default FEC (Ethernet 1/9 through 1/16)	Network Module Default FEC
25G-SR	Clause 74 FC-FEC	Clause 108 RS-FEC
25G-LR	Clause 74 FC-FEC	Clause 108 RS-FEC
10/25G-CSR	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-AOC *x* M	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-CU2.5/3M	Auto-Negotiate	Auto-Negotiate
25G-CU4/5M	Auto-Negotiate	Auto-Negotiate
25/50/100G	Clause 91 RS-FEC	Clause 91 RS-FEC

Note
From Version 7.4.2, Firewall Threat Defense supports the SFP 10G_25G_CSR_S module. To avoid link-down issues after a reboot, Forward Error Correction (FEC) must be enabled on these interfaces.

Step 6 (Optional) Refer to Configure IPv6 addressing, on page 1090 to configure IPv6 addressing on the **IPv6** tab.

Step 7 (Optional) Refer to Configure the MAC address, on page 1110 to manually configure the MAC address on the **Advanced** tab.

Step 8 Click **OK**.

Step 9 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure the bridge virtual interface (BVI)

Configure a Bridge Virtual Interface (BVI) to provide IP addressing for each bridge group, which enables traffic routing and management functionality for the bridge group members.

Each bridge group requires a BVI for which you configure an IP address. The Firewall Threat Defense uses this IP address as the source address for packets originating from the bridge group. The BVI IP address must

be on the same subnet as the connected network. For IPv4 traffic, the BVI IP address is required to pass any traffic. For IPv6 traffic, you must, at a minimum, configure the link-local addresses to pass traffic, but a global management address is recommended for full functionality, including remote management and other management operations.

For routed mode, if you provide a name for the BVI, then the BVI participates in routing. Without a name, the bridge group remains isolated as in transparent firewall mode.

Before you begin

You cannot add the BVI to a security zone; therefore, you cannot apply Access Control policies to the BVI. You must apply your policy to the bridge group member interfaces based on their zones.

Follow these steps to configure the Bridge Virtual Interface (BVI):

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** () for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Choose **Add Interfaces** > **Bridge Group Interface**.

a) (Routed Mode) In the **Name** field, enter a name up to 48 characters in length.

You must name the BVI if you want to route traffic outside the bridge group members, for example, to the outside interface or to members of other bridge groups. The name is not case-sensitive.

Step 3 In the **Bridge Group ID** field, enter the bridge group ID between 1 and 250.

Step 4 In the **Description** field, enter a description for this bridge group.

Step 5 On the **Interfaces** tab, click an interface and then click **Add** to move it to the **Selected Interfaces** area. Repeat for all interfaces that you want to make members of the bridge group.

Step 6 Click the **IPv4** tab.

a) (Routed Mode) To set the IP address, use one of the following options from the **IP Type** drop-down list.

High Availability and clustering interfaces only support static IP address configuration; DHCP is not supported.

- **Use Static IP**—Enter the IP address and subnet mask. For High Availability, you can only use a static IP address. Set the standby IP address on the **Devices** > **Device Management**, click the **High Availability** tab in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.
  - **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.

b) (Transparent Mode) In the **IP Address** field, enter the IPv4 address and subnet mask.

Do not assign a host address (/32 or 255.255.255.255) to the BVI. Also, do not use other subnets that contain fewer than 3 host addresses (one each for the upstream router, downstream router, and transparent firewall) such as a /30 subnet (255.255.255.252). The Firewall Threat Defense device drops all ARP

packets to or from the first and last addresses in a subnet. For example, if you use a /30 subnet and assign a reserved address from that subnet to the upstream router, then the Firewall Threat Defense device drops the ARP request from the downstream router to the upstream router.

For High Availability, set the standby IP address on the **Devices** > **Device Management** page, click the **High Availability** tab in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.

Step 7 (Optional) Refer to Configure IPv6 addressing, on page 1090 to configure IPv6 addressing.

Step 8 (Optional) Refer to Add a static ARP entry, on page 1111 and Add a static MAC address and disable MAC learning for a bridge group, on page 1112 (for transparent mode only) to configure the **ARP** and **MAC** settings.

Step 9 Click **OK**.

Click **Save**.

What to do next

Go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure IPv6 addressing

This section describes how to configure IPv6 addressing in routed and transparent mode.

About IPv6

This section includes information about IPv6.

IPv6 addressing

IPv6 addressing provides 128 bits of network addressing, which provides more than enough globally unique IP addresses. It supports both public global addresses and private link-local addresses, and

IPv6 address types

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network. For a bridge group, this address needs to be configured for the BVI, and not per member interface. You can also configure a global IPv6 address for the management interface in transparent mode.
- Link-local—The link-local address is a private address that you can only use on the directly-connected network. Routers do not forward packets using link-local addresses; they are only for communication on a particular physical network segment. They can be used for address configuration or for the Neighbor Discovery functions such as address resolution. In a bridge group, only member interfaces have link-local addresses; the BVI does not have a link-local address.

At a minimum, you need to configure a link-local address for IPv6 to operate. If you configure a global address, a link-local address is automatically configured on the interface, so you do not also need to specifically configure a link-local address. For bridge group member interfaces, when you configure the global address on the BVI, the Firewall Threat Defense device automatically generates link-local addresses for member

interfaces. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

Modified EUI-64 interface IDs

RFC 3513: Internet Protocol Version 6 (IPv6) Addressing Architecture requires that the interface identifier portion of all unicast IPv6 addresses, except those that start with binary value 000, be 64 bits long and be constructed in Modified EUI-64 format. The Firewall Threat Defense device can enforce this requirement for hosts attached to the local link.

Address verification process

When this feature is enabled on an interface, the source addresses of IPv6 packets received on that interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the IPv6 packets do not use the Modified EUI-64 format for the interface identifier, the packets are dropped and this system log message is generated:

```
325003: EUI-64 source address check failed.
```

The address format verification is only performed when a flow is created. Packets from an existing flow are not checked. Additionally, the address verification can only be performed for hosts on the local link.

IPv6 prefix delegation client configuration

The Firewall Threat Defense can act as a DHPCv6 Prefix Delegation client so that the client interface, for example the outside interface connected to a cable modem, can receive one or more IPv6 prefixes that the Firewall Threat Defense can then subnet and assign to its inside interfaces.

IPv6 prefix delegation

IPv6 prefix delegation enables the Firewall Threat Defense to function as a Dynamic Host Control Protocol for IPv6 (DHCPv6) prefix delegation client and receive one or more IPv6 prefixes. This feature allows the Firewall Threat Defense to subnet and assign these prefixes to its inside interfaces. It enable the hosts connected to inside interfaces to use Stateless Address Auto Configuration (SLAAC) to obtain global IPv6 addresses.

Client and server behavior

The client interface, such as the outside interface connected to a cable modem, receives IPv6 prefixes that the Firewall Threat Defense can then subnet and assign to its inside interfaces. The inside Firewall Threat Defense interfaces do not in turn act as Prefix Delegation servers; the Firewall Threat Defense can only provide global IP addresses to SLAAC clients.

If a router is connected to the Firewall Threat Defense, it can act as a SLAAC client to obtain its IP address. However, if you want to use a subnet of the delegated prefix for the networks behind the router, you must manually configure those addresses on the router's inside interfaces.

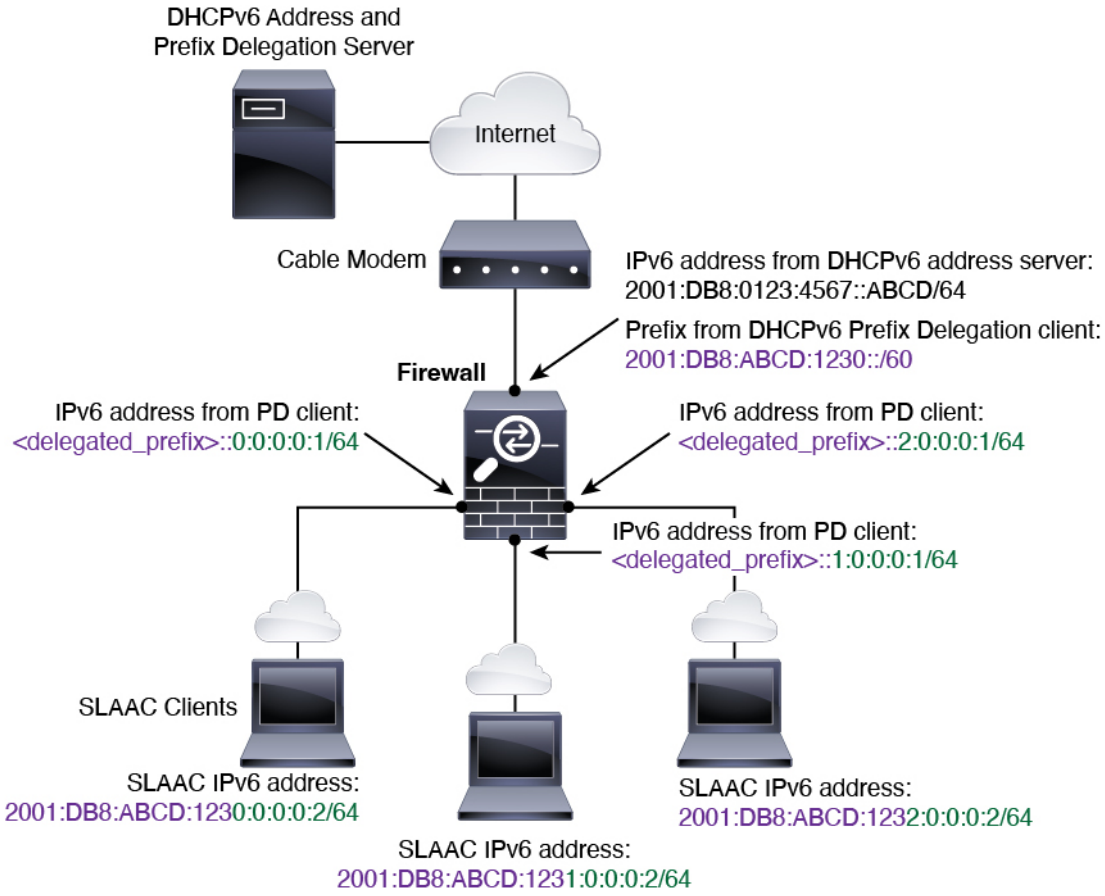
The Firewall Threat Defense includes a light DHCPv6 server. This server provides information, such as the DNS server and domain name, to SLAAC clients when they send Information Request (IR) packets to the Firewall Threat Defense. The Firewall Threat Defense only accepts IR packets, and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses based on prefixes received in Router Advertisement messages. In other words, the addresses are based on the prefix that the Firewall Threat Defense received using Prefix Delegation.

IPv6 prefix delegation /64 subnet example

Here is an IPv6 prefix delegation /64 subnet example that demonstrates how a device subnets delegated prefixes into /64 networks and assigns global IPv6 addresses to interfaces.

IPv6 prefix delegation /64 subnet configuration

This example shows the Firewall Threat Defense receiving an IP address on the outside interface using the DHCPv6 address client. It also gets a delegated prefix using the DHCPv6 Prefix Delegation client. The Firewall Threat Defense subnets the delegated prefix into /64 networks and assigns global IPv6 addresses to its inside interfaces dynamically using the delegated prefix plus a manually configured subnet (::0, ::1, or ::2) and IPv6 address (0:0:0:1) per interface. SLAAC clients connected to those inside interfaces obtain IPv6 addresses on each /64 subnet.

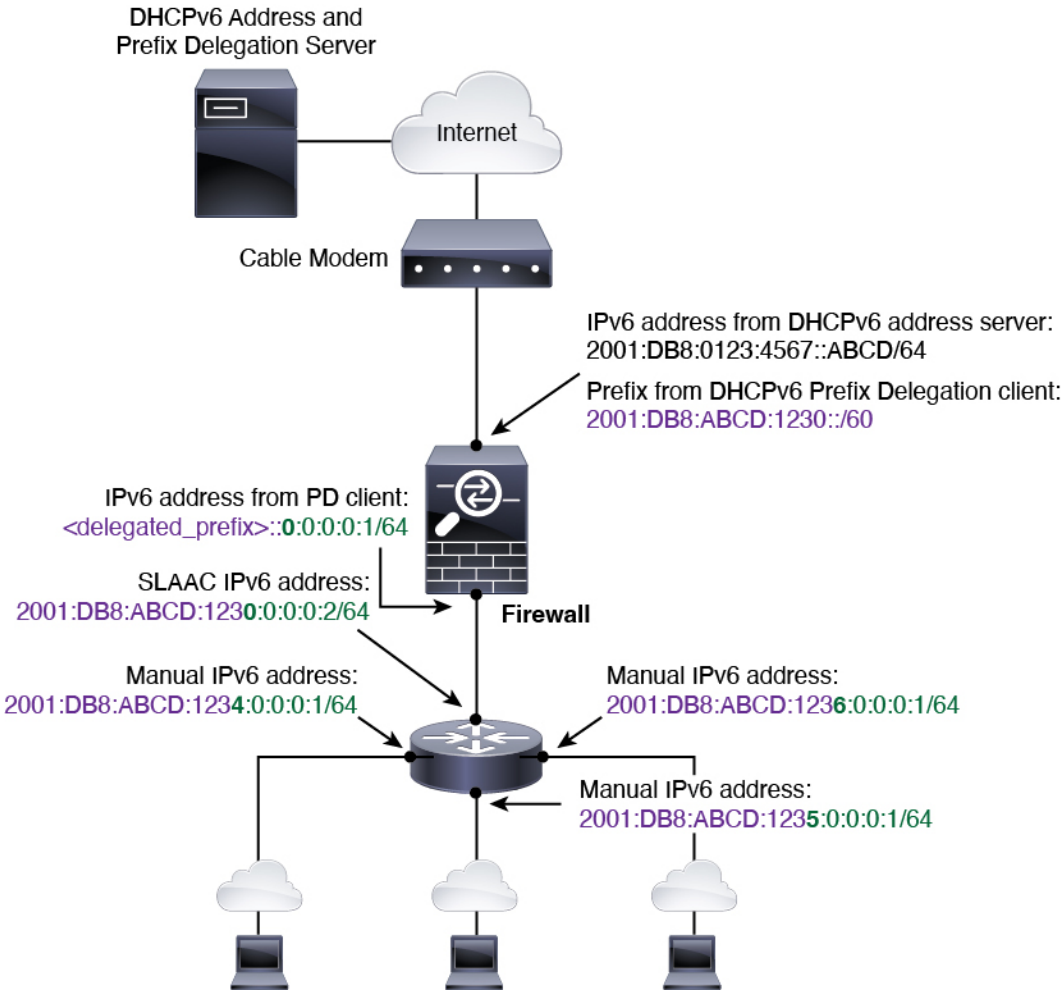


IPv6 prefix delegation /62 subnet examples

IPv6 prefix delegation /62 subnet enables subnetting of prefixes into 4 /62 subnets for downstream router delegation. It allows manual configuration of /64 subnets on router interfaces when dynamic prefix assignment is not available.

IPv6 /62 subnet delegation configuration

This example shows the Firewall Threat Defense subnetting the prefix into 4 /62 subnets: 2001:DB8:ABCD:123**0**::/62, 2001:DB8:ABCD:123**4**::/62, 2001:DB8:ABCD:123**8**::/62, and 2001:DB8:ABCD:123**C**::/62. The Firewall Threat Defense uses one of 4 available /64 subnets on 2001:DB8:ABCD:1230::/62 for its inside network (::0). You can then manually use additional /62 subnets for downstream routers. The router shown uses 3 of 4 available /64 subnets on 2001:DB8:ABCD:1234::/62 for its inside interfaces (::4, ::5, and ::6). In this case, the inside router interfaces cannot dynamically obtain the delegated prefix, so you need to view the delegated prefix on the Firewall Threat Defense, and then use that prefix for your router configuration. Usually, ISPs delegate the same prefix to a given client when the lease expires, but if the Firewall Threat Defense receives a new prefix, you will have to modify the router configuration to use the new prefix. The DHCP unique identifier (DUID) is persistent across reboots.



Enable the IPv6 prefix delegation client

Enable the DHCPv6 Prefix Delegation client to obtain one or more IPv6 prefixes that can be subnetted and assigned to inside networks.

Enable the DHCPv6 Prefix Delegation client on one or more interfaces. The Firewall Threat Defense obtains one or more IPv6 prefixes that it can subnet and assign to inside networks. Typically, the interface on which you enable the prefix delegation client obtains its IP address using the DHCPv6 address client; only other Firewall Threat Defense interfaces use addresses derived from the delegated prefix.

This feature is supported only in routed mode.

This feature is not supported in clustering or High Availability.

Before you begin

When you use Prefix Delegation, you must set the Firewall Threat Defense IPv6 neighbor discovery router advertisement interval to be much lower than the preferred lifetime of the prefix assigned by the DHCPv6 Server to prevent IPv6 traffic interruption. For example, if the DHCPv6 server sets the preferred Prefix Delegation lifetime to 300 seconds, you should set the Firewall Threat Defense RA interval to be 150 seconds. To set the preferred lifetime, use the **show ipv6 general-prefix** command. To set the Firewall Threat Defense RA interval, see Configure IPv6 neighbor discovery, on page 1099. The default is 200 seconds.

Procedure

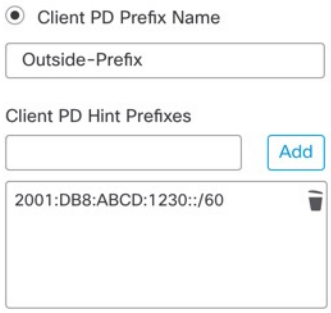
Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **IPv6** page, and then click **DHCP**.

Step 4 Click **Client PD Prefix Name** and enter a name for this prefix.

Figure 260: Enable the prefix delegation client



The name can be up to 200 characters.

Step 5 (Optional) Enter the prefix and prefix length in the **Client PD Hint Prefixes** field to provide one or more hints to the DHCP server about the delegated prefix you want to receive, then click **Add**.

Typically you want to request a particular prefix length, such as ::/60, or if you have received a particular prefix before and want to ensure you get it again when the lease expires, you can enter the whole prefix as the hint. If you enter multiple hints (different prefixes or lengths), then it is up to the DHCP server which hint to honor, or whether to honor the hint at all.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure a global IPv6 address

Before you begin

For IPv6 neighbor discovery for bridge groups, you must explicitly allow Neighbor Solicitation (ICMPv6 type 135) and Neighbor Advertisement (ICMPv6 type 136) packets through the Firewall Threat Defense bridge group member interfaces using a bidirectional access rule.

To configure a global IPv6 address for any routed mode interface and for the transparent or routed mode BVI, perform the following steps.

Note Configuring the global address automatically configures the link-local address, so you do not need to configure it separately. For bridge groups, configuring the global address on the BVI automatically configures link-local addresses on all member interfaces.

For subinterfaces defined on the Firewall Threat Defense , we recommend that you also set the MAC address manually, because they use the same burned-in MAC address of the parent interface. IPv6 link-local addresses are generated based on the MAC address, so assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the Firewall Threat Defense . See Configure the MAC address, on page 1110 .

Firewall Threat Defense Feature History:

- 7.3—DHCPv6 Address client
- 7.3—DHCPv6 Prefix Delegation client
- 7.3—DHCPv6 stateless server

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **IPv6** page.

For routed mode, the **Basic** page is selected by default. For transparent mode, the **Address** page is selected by default.

Step 4 On the **Basic** page, check **Enable IPv6** .

Use this option if you want to only configure the link-local addresses. Otherwise, configuring an IPv6 address enables IPv6 processing automatically.

Step 5 Configure the global IPv6 address using one of the following methods.

For failover and clustering, and for loopback interfaces, you must set the IP address manually. For clustering, manually configuring the link-local address is also not supported.

- (Routed interface) Stateless autoconfiguration—Check the **Autoconfiguration** check box.

  Enabling stateless autoconfiguration on the interface configures IPv6 addresses based upon prefixes received in Router Advertisement messages. A link-local address, based on the Modified EUI-64 interface ID, is automatically generated for the interface when stateless autoconfiguration is enabled.

  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the Firewall Threat Defense device does send Router Advertisement messages in this case. Click the **Settings** tab and click the **IPv6** tab and then uncheck the **Enable RA** check box to suppress messages.

- Manual configuration—To manually configure a global IPv6 address:

  a. Click the **Address** page, and click (+) **Add Address** .

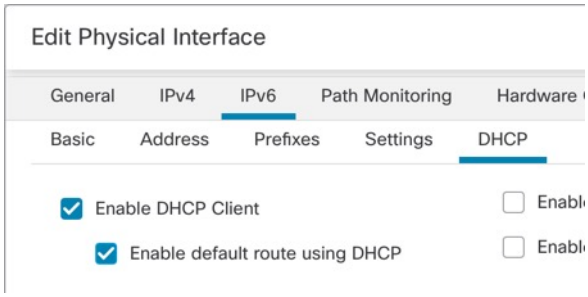
  The **Add Address** dialog box appears.

  b. In the **Address** field, enter either a full global IPv6 address, including the interface ID, or enter the IPv6 prefix, along with the IPv6 prefix length. (Routed Mode) If you only enter the prefix, then be sure to check the **Enforce EUI 64** check box to generate the interface ID using the Modified EUI-64 format. For example, 2001:0DB8::BA98:0:3210/48 (full address) or 2001:0DB8::/48 (prefix, with EUI 64 checked).

  For High Availability (if you did not set **Enforce EUI 64**), set the standby IP address on the **Devices** > **Device Management** page, click **High Availability** in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.

- (Routed interface) Obtain an address using DHCPv6—To use DHCPv6:

  Figure 261: Enable the DHCPv6 Client



  a. Click the **DHCP** page.

  b. Check the check box of **Enable DHCP Client** .

  c. (Optional) Check the check box of **Enable default route using DHCP** to obtain a default route from Router Advertisements.

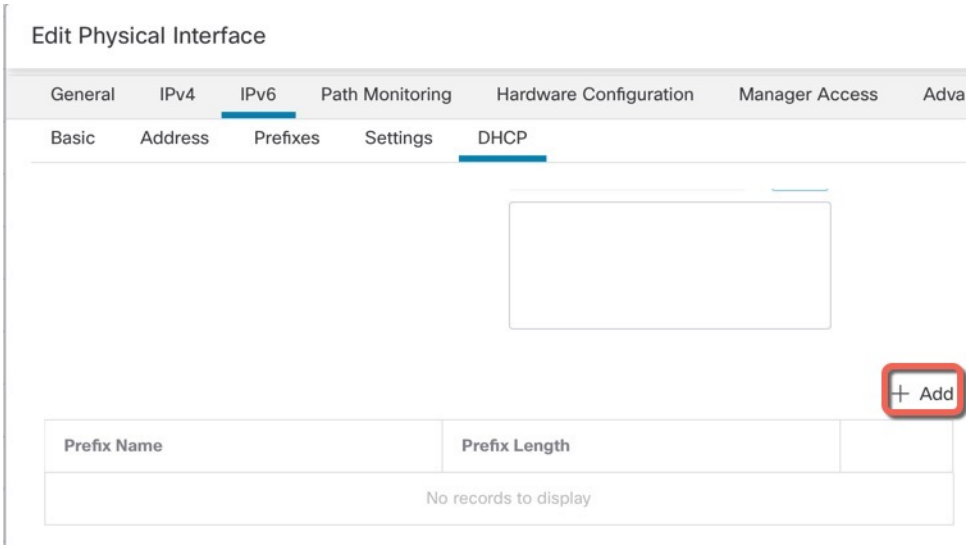
- (Routed interface) Use a delegated prefix—To assign an IPv6 address using the delegated prefix:

  This feature requires the Firewall Threat Defense to have the DHCPv6 Prefix Delegation client enabled *on a different interface* . See Enable the IPv6 prefix delegation client, on page 1093 .

a. Click the **DHCP** page.

b. Click (+) .

Figure 262: Use a Delegated Prefix

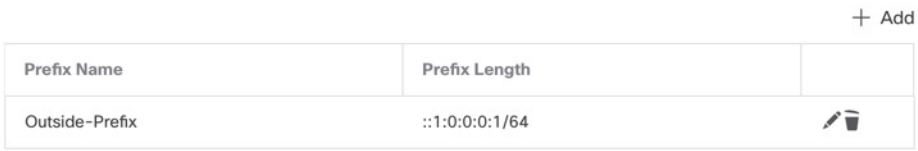


c. Enter the IPv6 address and **Prefix Length** .

Typically, the delegated prefix will be /60 or smaller so you can subnet to multiple /64 networks. /64 is the supported subnet length if you want to support SLAAC for connected clients. You should specify an address that completes the /60 subnet, for example ::1:0:0:0:1. Enter :: before the address in case the prefix is smaller than /60. For example, if the delegated prefix is 2001:DB8:1234:5670::/60, then the global IP address assigned to this interface is 2001:DB8:1234:5671::1/64. The prefix that is advertised in router advertisements is 2001:DB8:1234:5671::/64. In this example, if the prefix is smaller than /60, the remaining bits of the prefix will be 0's as indicated by the leading ::. For example, if the prefix is 2001:DB8:1234::/48, then the IPv6 address will be 2001:DB8:1234::1:0:0:0:1/64.

d. Click **OK** .

Figure 263: Prefix Delegation Table



e. Optionally enable the DHCPv6 stateless server on this interface (see Enable the DHCPv6 stateless server, on page 1232). If you do so, we recommend that you also check the **Enable DHCP for non-address config** option.

Step 6 For Routed interfaces, you can optionally set these values on the **Basic** page:

- To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.

- To manually set the link-local address, enter an address in the **Link-Local address** field.

  A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. If you do not want to configure a global address, and only need to configure a link-local address, you have the option of manually defining the link-local address. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

  Clustering does not support manual link-local addresses.

Step 7 For Routed interfaces, you can optionally set the following values on the **DHCP** page:

- Check the **Enable DHCP for address config** check box to set the Managed Address Config flag in the IPv6 router advertisement packet.

  This flag in IPv6 router advertisements informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.

- Check the **Enable DHCP for non-address config** check box to set the Other Address Config flag in the IPv6 router advertisement packet.

  This flag in IPv6 router advertisements informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address. Use this option when using the DHCPv6 stateless server with DHCPv6 prefix delegation.

Step 8 For Routed interfaces, see Configure IPv6 neighbor discovery, on page 1099 to configure settings on the **Prefixes** and **Settings** pages. For BVI interfaces, see the following parameters on the **Settings** page:

- **DAD attempts** —The maximum number of DAD attempts , between 1 and 600. Set the value to 0 to disable duplicate address detection (DAD) processing. This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses . 1 attempt is the default.
- **NS Interval** —The interval between IPv6 neighbor solicitation retransmissions on an interface , between 1000 and 3600000 ms . The default value is 1000 ms.
- **Reachable Time** —The amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred , between 0 and 3600000 ms . The default value is 0 ms. When 0 is used for the value, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value. The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly, however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

Step 9 Click **OK** .

Click **Save**.

What to do next

Go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure IPv6 neighbor discovery

Before you begin

Supported in Routed mode only. For IPv6 neighbor settings supported in transparent mode, see Configure a global IPv6 address, on page 1095 .

The IPv6 neighbor discovery process uses ICMPv6 messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the readability of a neighbor, and keep track of neighboring routers.

Nodes (hosts) use neighbor discovery to determine the link-layer addresses for neighbors known to reside on attached links and to quickly purge cached values that become invalid. Hosts also use neighbor discovery to find neighboring routers that are willing to forward packets on their behalf. In addition, nodes use the protocol to actively keep track of which neighbors are reachable and which are not, and to detect changed link-layer addresses. When a router or the path to a router fails, a host actively searches for functioning alternates.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click **IPv6** , and then **Prefixes**.

To configure the IPv6 router advertisements with IPv6 prefixes, perform these steps:

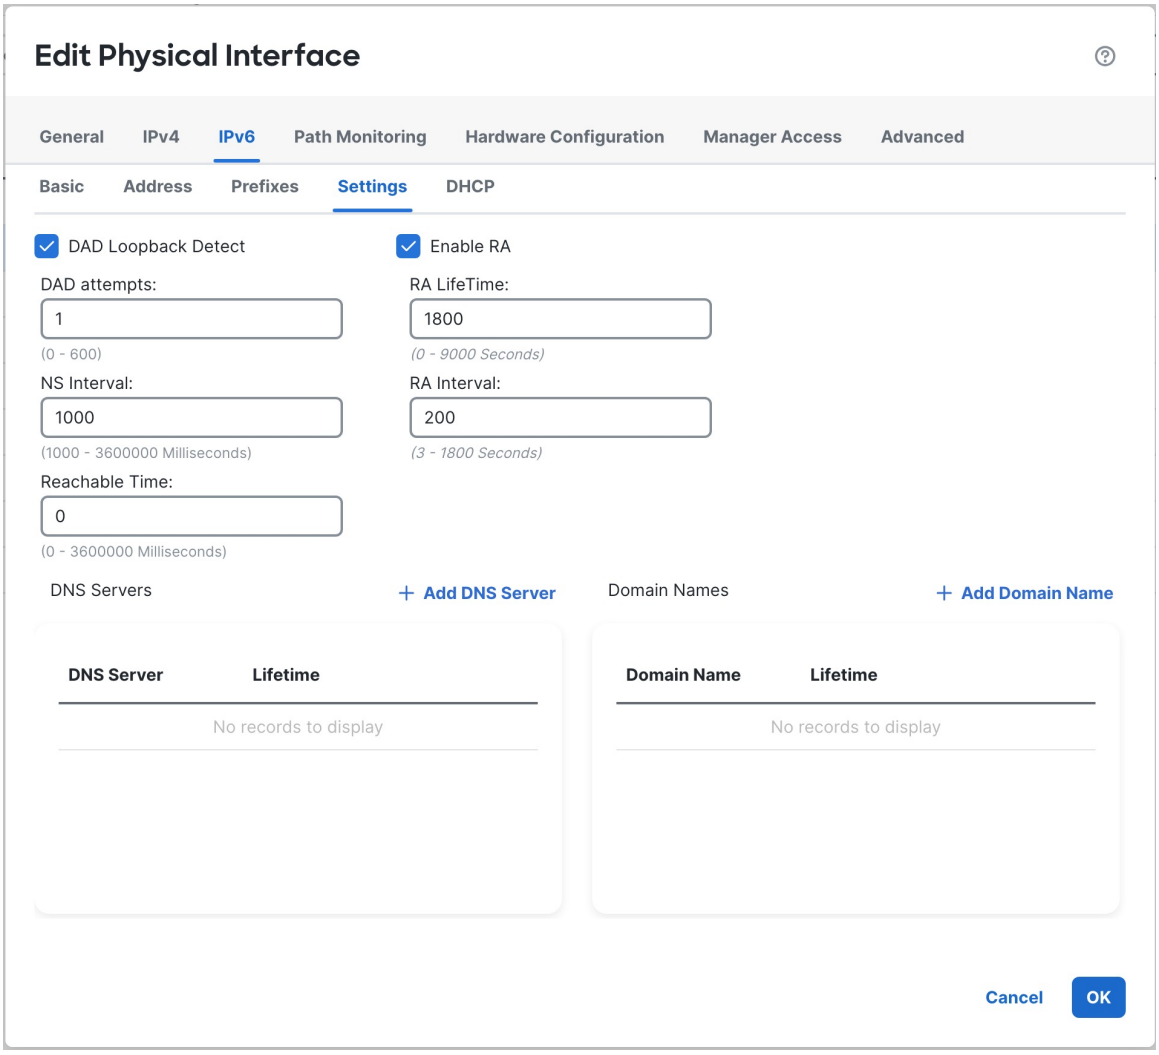
a) Click (+) **Add Prefix** .
b) In the **Address** field, enter the IPv6 address with the prefix length or check the **Default** check box to use the default prefix.
c) (Optional) Uncheck the **Advertisement** check box to indicate that the IPv6 prefix is not advertised. For the **Default** prefix, this setting only applies to on-link prefixes. Off-link prefixes will still be advertised unless you uncheck **Advertisement** for a specific off-link prefix.
d) Check the **Off Link** check box to indicate that the specified prefix is assigned to the link. Nodes sending traffic to addresses that contain the specified prefix consider the destination to be locally reachable on the link. This prefix should not be used for on-link determination.
e) To use the specified prefix for autoconfiguration, check the **Autoconfiguration** check box.
f) For the **Prefix Lifetime** , click **Duration** or **Expiration Date** .

- **Duration** —Enter a **Preferred Lifetime** for the prefix in seconds. This setting is the amount of time that the specified IPv6 prefix is advertised as being valid. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default is 2592000 (30 days). Enter a **Valid Lifetime** for the prefix in seconds. This setting is the amount of time that the specified IPv6 prefix is advertised as being preferred. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default setting is 604800 (seven days). Alternatively, check the **Infinite** check box to set an unlimited duration.
- **Expiration Date** —Choose a **Valid** and **Preferred** date and time.

g) Click **OK** .

Step 4 Click **Settings**.

Figure 264: IPv6 Settings



a) (Optional) Set the maximum number of **DAD attempts** , between 1 and 600. 1 attempt is the default. Set the value to 0 to disable duplicate address detection (DAD) processing.

   This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses.

   During the stateless autoconfiguration process, Duplicate Address Detection verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces.

   When a duplicate address is identified, the state of the address is set to DUPLICATE, the address is not used, and the following error message is generated:

   ```
   325002: Duplicate address ipv6_address/MAC_address on interface
   ```

   If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface. If the duplicate address is a global address, the address is not used.

b) (Optional) Configure the interval between IPv6 neighbor solicitation retransmissions in the **NS Interval** field, between 1000 and 3600000 ms.

The default value is 1000 ms.

Neighbor solicitation messages (ICMPv6 Type 135) are sent on the local link by nodes attempting to discover the link-layer addresses of other nodes on the local link. After receiving a neighbor solicitation message, the destination node replies by sending a neighbor advertisement message (ICPMv6 Type 136) on the local link.

After the source node receives the neighbor advertisement, the source node and destination node can communicate. Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verifying the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link.

Step 5 (Optional) Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred in the **Reachable Time** field, between 0 and 3600000 ms.

The default value is 0 ms. When 0 is used for the value, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value.

The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly, however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

Step 6 (Optional) To suppress the router advertisement transmissions, uncheck the **Enable RA** check box. If you enable router advertisement transmissions, you can set the RA lifetime and interval.

Router advertisement messages (ICMPv6 Type 134) are automatically sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message.

You may want to disable these messages on any interface for which you do not want the Firewall Threat Defense to supply the IPv6 prefix (for example, the outside interface).

- **RA Lifetime** —Configure the router lifetime value in IPv6 router advertisements, between 0 and 9000 seconds.

  The default is 1800 seconds.

- **RA Interval** —Configure the interval between IPv6 router advertisement transmissions, between 3 and 1800 seconds.

  The default is 200 seconds.

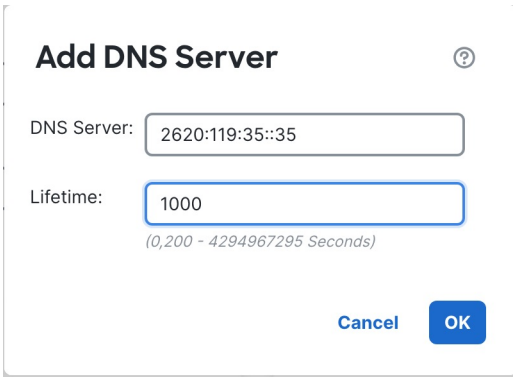
  To prevent synchronization with other IPv6 nodes, the firewall randomly adjusts the value that you set (jitter).

Step 7 (Optional) Configure the Recursive DNS Server (RDNSS) option to advertise a list of DNS servers to IPv6 clients.

This option is useful for clients that use SLAAC instead of DHCPv6 for their addressing. Note that if you enabled the prefix delegation client on the Firewall Threat Defense , you can alternatively pass along the DNS server information that the Firewall Threat Defense received using the Firewall Threat Defense 's DHCPv6 stateless server. If you configure both methods, the client will receive both sets of servers.

a) In the **DNS Servers** area, click **Add DNS Server** .

Figure 265: Add a DNS Server



b) Enter the **DNS Server** IPv6 Address.

c) Enter the maximum **Lifetime** , in seconds, that the DNS server will be used for name resolution, between 200 and 4294967295. If you set the value to 0, the entry will not be used. If you set the value to 4294967295, the entry will never expire. The default is 3 x the maximum RA interval. The value must be greater than or equal to the **RA Interval** .

d) Click **OK** .

e) Repeat for up to 8 servers.

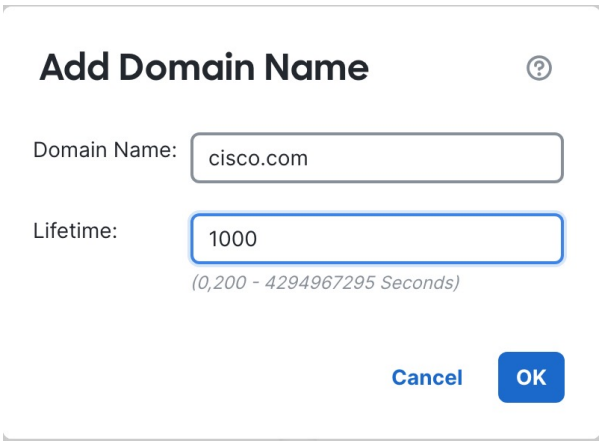
The servers are advertised in the order you add them.

Step 8 (Optional) Configure the DNS Search List (DNSSL) option to advertise a list of search domains to IPv6 clients.

This option is useful for clients that use SLAAC instead of DHCPv6 for their addressing. Note that if you enabled the prefix delegation client on the Firewall Threat Defense , you can alternatively pass along the DNS domain that the Firewall Threat Defense received using the Firewall Threat Defense 's DHCPv6 stateless server. If you configure both methods, the client will receive both sets of domains.

a) In the **Router Advertisement DNS-SEARCH-LIST** area, click **Add** .

Figure 266: Add a Domain for Search



b) Enter the **Domain Name** .

c) Enter the maximum **Lifetime** , in seconds, that the domain will be used, between 200 and 4294967295. If you set the value to 0, the entry will not be used. If you set the value to 4294967295, the entry will never expire. The default is 3 x the maximum RA interval. The value must be greater than or equal to the **RA Interval** .
d) Click **OK** .
e) Repeat for up to 5 domains.

The domains are advertised in the order you add them.

Step 9 Click **OK**.

Click **Save**.

What to do next

Go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Advanced interface settings

This section describes how to configure MAC addresses for regular firewall mode interfaces, how to set the maximum transmission unit (MTU), and how to set other advanced parameters.

About advanced interface configuration

This section describes advanced interface settings.

MAC addresses

Media Access Control (MAC) address is a unique physical address of a network device.

You can manually assign MAC addresses to override the default. For container instances, the FXOS chassis automatically generates unique MAC addresses for all interfaces.

Note You might want to assign unique MAC addresses to subinterfaces defined on the Firewall Threat Defense, because they use the same burned-in MAC address of the parent interface. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the Firewall Threat Defense device.

Note For container instances, even if you are not sharing a subinterface, if you manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification.

Default MAC addresses

This reference provides information about default MAC addresses used in network devices and their configurations.

Default MAC address assignments depend on the type of interface.

- Physical interfaces—The physical interface uses the burned-in MAC address.
- VLAN interfaces (Firepower 1010 and Secure Firewall 1210/1220)—Routed firewall mode: All VLAN interfaces share a MAC address. Ensure that any connected switches can support this scenario. If the connected switches require unique MAC addresses, you can manually assign MAC addresses. See Configure the MAC address, on page 1110.

  Transparent firewall mode: Each VLAN interface has a unique MAC address. You can override the generated MAC addresses if desired by manually assigning MAC addresses. See Configure the MAC address, on page 1110.
- EtherChannels—For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses a unique MAC address from a pool; interface membership does not affect the MAC address.
- EtherChannels (ASA Models)—The port-channel interface uses the lowest-numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can configure a MAC address for the port-channel interface. We recommend configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.
- Subinterfaces (Firewall Threat Defense-defined)—All subinterfaces of a physical interface use the same burned-in MAC address. You might want to assign unique MAC addresses to subinterfaces. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the Firewall Threat Defense.

For container instances:

- MAC addresses for all interfaces are taken from a MAC address pool. For subinterfaces, if you decide to manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification. See Automatic MAC Addresses for Container Instance Interfaces.

MTU

Maximum Transmission Unit (MTU) specifies the maximum frame *payload* size that the Firewall Threat Defense can transmit on a given Ethernet interface. The MTU value *excludes* Ethernet headers, VLAN tags, or other overhead. For example, if you set the MTU to 1500, the expected frame size is 1518 bytes including the headers, or 1522 when VLAN tagging is used. Do not set the MTU value higher to accommodate these headers.

VXLAN and geneve encapsulation requirements

For Geneve, the entire Ethernet datagram is being encapsulated, so the new IP packet is larger and requires a larger MTU. Set the Firewall Threat Defense VTEP source interface MTU to the network MTU plus 306 bytes.

Path MTU discovery

Path MTU Discovery enables all devices in a network path between two hosts to coordinate the MTU so they can standardize on the lowest MTU in the path.

Path MTU discovery support

The Firewall Threat Defense supports Path MTU Discovery (as defined in RFC 1191).

Default MTU

The default MTU on the Firewall Threat Defense device is 1500 bytes. This value does not include the 18-22 bytes for the Ethernet header, VLAN tagging, or other overhead.

MTU and fragmentation

IPv4 packets that are larger than the specified MTU are fragmented into multiple frames.

Packet handling details

For IPv4, if an outgoing IP packet is larger than the specified MTU, it is fragmented into 2 or more frames. Fragments are reassembled at the destination (and sometimes at intermediate hops), and fragmentation can cause performance degradation. For IPv6, packets are typically not allowed to be fragmented at all. Therefore, your IP packets should fit within the MTU size to avoid fragmentation.

For TCP packets, the endpoints typically use their MTU to determine the TCP maximum segment size (MTU - 40, for example). If additional TCP headers are added along the way, for example for site-to-site VPN tunnels, then the TCP MSS might need to be adjusted down by the tunneling entity. See TCP MSS, on page 1106.

For UDP or ICMP, the application should take the MTU into account to avoid fragmentation.



Note The Firewall Threat Defense can receive frames larger than the configured MTU as long as there is room in memory.

Full reassembly for inspected fragmented packets

The FTD performs virtual reassembly by default. However, when features such as DNS or voice media inspection are active, the firewall performs full reassembly before forwarding. This ensures packets are processed correctly by inspection engines.

For example, SIP packets subject to inspection require full reassembly. If SIP inspection is disabled, the system maintains virtual reassembly. To avoid issues, we recommend reducing SIP message size or using TCP or TLS for large messages rather than sending large SIP INVITE packets over UDP.

The firewall merges fragments in memory. If the egress interface MTU is smaller than the reassembled packet, the device will refragment the data to fit. To verify behavior, use packet captures on the egress interface and adjust MTU settings as needed.

MTU and jumbo frames configuration

A larger MTU lets you send larger packets. Larger packets might be more efficient for your network.

- Matching MTUs on the traffic path—We recommend that you set the MTU on all Firewall Threat Defense interfaces and other device interfaces along the traffic path to be the same. Matching MTUs prevents intermediate devices from fragmenting the packets.
- Accommodating jumbo frames—You can set the MTU 9000 bytes or higher when you enable jumbo frames. The maximum depends on the model.

TCP MSS

The TCP maximum segment size (MSS) is the size of the TCP payload *before* any TCP and IP headers are added. TCP MSS values are exchanged between client and server during the three-way handshake when establishing connections. UDP packets are not affected.

TCP MSS configuration and behavior

You can set the TCP MSS on the Firewall Threat Defense device for through traffic using the Sysopt_Basic object in FlexConfig. By default, the maximum TCP MSS is set to 1380 bytes. This setting is useful if the Firewall Threat Defense device must increase the packet size for IPsec VPN encapsulation. However, for non-IPsec endpoints, you should disable the maximum TCP MSS on the Firewall Threat Defense device.

If you set a maximum TCP MSS and either endpoint of a connection requests a TCP MSS larger than this value, the Firewall Threat Defense device overwrites the TCP MSS in the request packet with the Firewall Threat Defense device maximum. If the host or server does not request a TCP MSS, the Firewall Threat Defense device assumes the RFC 793-default value of 536 bytes for IPv4 or 1220 bytes for IPv6, but does not modify the packet.

You can also configure the minimum TCP MSS. If a host or server requests a very small TCP MSS, the Firewall Threat Defense device can increase the value. By default, the minimum TCP MSS is not enabled.

For to-the-box traffic, including for SSL VPN connections, this setting does not apply. The Firewall Threat Defense device derives the TCP MSS from MTU: MTU minus 40 for IPv4 or MTU minus 60 for IPv6.

TCP MSS adjustment scenario

For example, you leave the default MTU as 1500 bytes. A host requests an MSS of 1500 minus the TCP and IP header length, which sets the MSS to 1460. If the Firewall Threat Defense device maximum TCP MSS is 1380 (the default), then the Firewall Threat Defense device changes the MSS value in the TCP request packet to 1380. The server then sends packets with 1380-byte payloads. The Firewall Threat Defense device can add up to 120 bytes of headers to the packet, and the resulting packet fits within the MTU size of 1500 bytes.

Default TCP MSS

A default TCP MSS is a network configuration setting that sets the maximum TCP MSS on the Firewall Threat Defense device to 1380 bytes by default, accommodates IPv4 IPsec VPN connections where the headers can equal up to 120 bytes, and fits within the default MTU of 1500 bytes.

Suggested maximum TCP MSS setting

The default TCP MSS assumes the Firewall Threat Defense device acts as an IPv4 IPsec VPN endpoint and has an MTU of 1500. When the Firewall Threat Defense device acts as an IPv4 IPsec VPN endpoint, it needs to accommodate up to 120 bytes for TCP and IP headers.

If you change the MTU value, use IPv6, or do not use the Firewall Threat Defense device as an IPsec VPN endpoint, then you should change the TCP MSS setting using the Sysopt_Basic object in FlexConfig.



Note If you explicitly set an MSS, a component such as TLS/SSL decryption or server discovery might require a particular MSS.In that case, the component sets the MSS based on the interface MTU and ignores your MSS setting.

- Normal traffic—Disable the TCP MSS limit and accept the value established between connection endpoints. Because connection endpoints typically derive the TCP MSS from the MTU, non-IPsec packets usually fit this TCP MSS.
- IPv4 IPsec endpoint traffic—Set the maximum TCP MSS to the MTU - 120. For example, if you use jumbo frames and set the MTU to 9000, then you need to set the TCP MSS to 8880 to take advantage of the new MTU.
- IPv6 IPsec endpoint traffic—Set the maximum TCP MSS to the MTU - 140.

ARP inspection for bridge group traffic

This reference describes ARP inspection functionality for bridge group traffic in transparent firewall configurations.

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

When you enable ARP inspection, the Firewall Threat Defense device compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the Firewall Threat Defense device drops the packet.
- If the ARP packet does not match any entries in the static ARP table, then you can set the Firewall Threat Defense device to either forward the packet out all interfaces (flood), or to drop the packet.

Note The dedicated Management interface never floods packets even if this parameter is set to flood.

MAC address tables

A MAC address table associates each MAC address with its source interface for accurate packet routing.

MAC address table behavior

When you use bridge groups, the Firewall Threat Defense learns and builds a MAC address table in a way similar to a normal bridge or switch. When a device sends a packet through the bridge group, the Firewall Threat Defense adds the MAC address to its table. The table associates the MAC address with the source interface so that the Firewall Threat Defense knows to send any packets addressed to the device out the correct interface. Because traffic between bridge group members is subject to the Firewall Threat Defense security policy, if the destination MAC address of a packet is not in the table, the Firewall Threat Defense does not flood the original packet on all interfaces as a normal bridge does. Instead, it generates the packets for directly-connected devices or for remote devices.

- Packets for directly-connected devices—The Firewall Threat Defense generates an ARP request for the destination IP address so that it can learn which interface receives the ARP response.
- Packets for remote devices—The Firewall Threat Defense generates a ping to the destination IP address so that it can learn which interface receives the ping reply.

The original packet is dropped.

Default settings

This reference provides default configuration settings for ARP inspection, MAC address table entries, and interface behavior for MAC address learning.

- If you enable ARP inspection, the default setting is to flood non-matching packets.
- The default timeout value for dynamic MAC address table entries is five minutes.
- By default, each interface automatically learns the MAC addresses of entering traffic, and the Firewall Threat Defense device adds corresponding entries to the MAC address table.

Note Secure Firewall Threat Defense device generates a reset packet to reset a connection that is denied by a stateful inspection engine. Here, the destination MAC address of the packet is not determined based on the ARP table lookup but instead it is taken directly from the packets (connections) that are being denied.

Guidelines for ARP inspection and the MAC address table

Consider these configuration support limitations for ARP inspection and MAC address table management:

- ARP inspection is only supported for bridge groups.
- MAC address table configuration is only supported for bridge groups.

Configure the MTU

Customize the MTU on the interface, for example, to allow jumbo frames.

For the ISA 3000 and the Firewall Threat Defense Virtual: Changing the MTU above 1500 bytes automatically enables jumbo-frame reservation. You must restart the system before you can use jumbo frames. For the Firewall Threat Defense Virtual that supports clustering, you can enable jumbo-frame reservation in the Day0 configuration, so in that case, you do not need to restart. After you restart, you cannot disable jumbo-frame reservation. An exception is for the Firewall Threat Defense Virtual, where you can disable jumbo-frame reservation in the Day0 configuration, if supported. If you use an interface in an inline set, the MTU setting is not used. However, the jumbo-frame reservation setting *is* relevant to inline sets; jumbo frames enable the inline interfaces to receive packets up to 9000 bytes. To enable jumbo-frame reservation, you must set the MTU of *any* interface above 1500 bytes.

Jumbo frames are enabled by default on other platforms.

Caution Changing the highest MTU value on the device for a data interface restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all data interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. This caution does not apply to management-only interfaces. See Snort restart traffic behavior, on page 196 for more information.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 On the **General** tab, set the **MTU**. The minimum and maximum depends on your platform.

The default is 1500 bytes.

Step 4 Click **OK**.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 6 For the ISA 3000 and the Firewall Threat Defense Virtual, if you set the MTU above 1500 bytes, restart the system to enable jumbo-frame reservation. See Shut down or restart the device, on page 107.

Configure the MAC address

This task allows you to manually assign MAC addresses to interfaces, including active and standby MAC addresses for high availability configurations.

You might need to manually assign a MAC address. You can also set the Active and Standby MAC addresses on the **Devices** > **Device Management** window by choosing **High Availability** from the **Add** drop-down list. If you set the MAC address for an interface on both screens, the addresses on the **Interfaces** > **Advanced** tab take precedence.

Note For container instances, even if you are not sharing a subinterface, if you manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Advanced** tab.

The **Information** tab is selected.

Step 4 (For clustering in individual interface mode) Choose a MAC address pool from the drop-down list.

You can add MAC address pools according to Configure address pools, on page 2954.

Step 5 (For other modes) Set the active and standby MAC addresses.

a) In the **Active MAC Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

b) In the **Standby MAC Address** field, enter a MAC address for use with High Availability.

If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add a static ARP entry

Add static ARP entries to control ARP packet flow and support ARP inspection for bridge group members or routed interfaces.

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection (see Configure ARP inspection, on page 1244). ARP inspection compares ARP packets with *static* ARP entries in the ARP table.

For routed interfaces, you can enter static ARP entries, but dynamic entries are usually sufficient. For routed interfaces, the ARP table is used to deliver packets to directly-connected hosts. Although senders identify a packet destination by an IP address, the actual delivery of the packet on Ethernet relies on the Ethernet MAC address. When a router or host wants to deliver a packet on a directly connected network, it sends an ARP request asking for the MAC address associated with the IP address, and then delivers the packet to the MAC address according to the ARP response. The host or router keeps an ARP table so it does not have to send ARP requests for every packet it needs to deliver. The ARP table is dynamically updated whenever ARP responses are sent on the network, and if an entry is not used for a period of time, it times out. If an entry is incorrect (for example, the MAC address changes for a given IP address), the entry must time out before it can be updated with the new information.

For transparent mode, the Firewall Threat Defense only uses dynamic ARP entries in the ARP table for traffic to and from the Firewall Threat Defense device, such as management traffic.

Before you begin

This screen is only available for named interfaces.

Follow these steps to add a static ARP entry:

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **ARP** tab (called **ARP and MAC** for transparent mode).

Step 4 Click (+)**Add ARP Config**.
The **Add ARP Config** dialog box appears.

Step 5 In the **IP Address** field, enter the IP address of the host.

Step 6 In the **MAC Address** field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b.

Step 7 To perform proxy ARP for this address, check the **Enable Alias** check box.

If the Firewall Threat Defense device receives an ARP request for the specified IP address, then it responds with the specified MAC address.

Step 8 Click **OK**, and then click **OK** again to exit the Advanced settings.

Step 9 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add a static MAC address and disable MAC learning for a bridge group

Add static MAC addresses to the MAC address table and optionally disable MAC address learning to guard against MAC spoofing and control traffic flow through the Firewall Threat Defense device.

Normally, MAC addresses are added to the MAC address table dynamically as traffic from a particular MAC address enters an interface. You can disable MAC address learning. If you do not statically add MAC addresses to the table, no traffic can pass through the Firewall Threat Defense device. You can also add static MAC addresses to the MAC address table. One benefit to adding static entries is to guard against MAC spoofing. If a client uses the same MAC address as a static entry and sends traffic to an interface that does not match the static entry, the Firewall Threat Defense device drops the traffic and generates a system message. When you add a static ARP entry (see Add a static ARP entry, on page 1111), a static MAC address entry is automatically added to the MAC address table.

Before you begin

This screen is only available for named BVIs in transparent mode.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **ARP and MAC** tab.

Step 4 (Optional) Disable MAC learning by unchecking the **Enable MAC Learning** check box.

Step 5 To add a static MAC address, click **Add MAC Config**.
The **Add MAC Config** dialog box appears.

Step 6 In the **MAC Address** field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b. Click **OK**.

Step 7 Click **OK** to exit the Advanced settings.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Set security configuration parameters

This task configures security parameters on named interfaces to prevent IP spoofing attacks and control packet fragmentation behavior. You can enable anti-spoofing protection, configure full fragment reassembly, and customize fragment handling settings.

This section describes how to prevent IP spoofing, allow full fragment reassembly, and override the default fragment setting at the device level in **Platform Settings**.

Anti-spoofing

This section lets you enable Unicast Reverse Path Forwarding on an interface. Unicast RPF guards against IP spoofing (a packet uses an incorrect source IP address to obscure its true source) by ensuring that all packets have a source IP address that matches the correct source interface according to the routing table.

Normally, the Firewall Threat Defense device only looks at the destination address when determining where to forward the packet. Unicast RPF instructs the device to also look at the source address. This is why it is called Reverse Path Forwarding. For any traffic that you want to allow through the Firewall Threat Defense device, the device routing table must include a route back to the source address. See RFC 2267 for more information.

For outside traffic, for example, the Firewall Threat Defense device can use the default route to satisfy the Unicast RPF protection. If traffic enters from an outside interface, and the source address is not known to the routing table, the device uses the default route to correctly identify the outside interface as the source interface.

If traffic enters the outside interface from an address that is known to the routing table, but is associated with the inside interface, then the Firewall Threat Defense device drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the device drops the packet because the matching route (the default route) indicates the outside interface.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

Fragment per packet

By default, the Firewall Threat Defense device allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments through the Firewall Threat Defense device. Fragmented packets are often used as DoS attacks.

Fragment reassembly

The Firewall Threat Defense device performs the following fragment reassembly processes:

- IP fragments are collected until a fragment set is formed or until a timeout interval has elapsed.
- If a fragment set is formed, integrity checks are performed on the set. These checks include no overlapping, no tail overflow, and no chain overflow.
- IP fragments that terminate at the Firewall Threat Defense device are always fully reassembled.
- If **Full Fragment Reassembly** is disabled (the default), the fragment set is forwarded to the transport layer for further processing.
- If **Full Fragment Reassembly** is enabled, the fragment set is first coalesced into a single IP packet. The single IP packet is then forwarded to the transport layer for further processing.

Full reassembly for inspected fragmented packets

The Firewall Threat Defense performs virtual reassembly by default. However, when features such as DNS or voice media inspection are active, the firewall performs full reassembly before forwarding. This ensures packets are processed correctly by inspection engines.

For example, SIP packets subject to inspection require full reassembly. If SIP inspection is disabled, the system maintains virtual reassembly. To avoid issues, we recommend reducing SIP message size or using TCP or TLS for large messages rather than sending large SIP INVITE packets over UDP.

The firewall merges fragments in memory. If the egress interface MTU is smaller than the reassembled packet, the device will refragment the data to fit. To verify behavior, use packet captures on the egress interface and adjust MTU settings as needed.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **Security Configuration** tab.

Step 4 To enable Unicast Reverse Path Forwarding, check the **Enable Anti Spoofing** check box.

Step 5 To enable full fragment reassembly, check the **Allow Full Fragment Reassembly** check box.

Step 6 To change the number of fragments allowed per packet, check the **Override Default Fragment Setting** check box, and set the following values:

- **Size**—Set the maximum number of packets that can be in the IP reassembly database waiting for reassembly. The default is 200. Set this value to 1 to disable fragments.
- **Chain**—Set the maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets.
- **Timeout**—Set the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded. The default is 5 seconds.

Step 7 Click **OK**.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

History for Regular Firewall Interfaces

Provides a comprehensive history of interface-related features, including hardware support, protocol enhancements, and configuration changes for Secure Firewall devices.

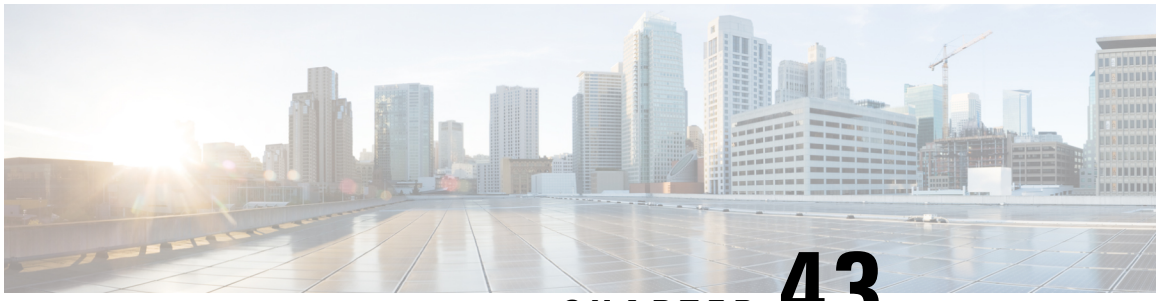
Table 69: History for Regular Firewall Interfaces

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Recursive DNS Server (RDNSS) and DNS Search List (DNSSL) options to advertise a list of DNS servers and domains to IPv6 clients	10.0.0 20251121	10.0.0	You can now configure Recursive DNS Server (RDNSS) and DNS Search List (DNSSL) options to provide DNS servers and domains to SLAAC clients using router advertisements. New/modified screens: **Devices** > **Device Management** > **Interfaces** > **Add/Edit Interfaces** > **IPv6** > **Settings** New/modified commands: **show ipv6 nd detail** , **show ipv6 nd ra dns-search-list** , **show ipv6 nd ra dns server** , **show ipv6 nd summary**
Secure Firewall 200 hardware switch support	10.0.0 20251121	10.0.0	The Secure Firewall 200 supports setting each Ethernet interface to be a switch port or a firewall interface.
Secure Firewall 1210CP IEEE 802.3bt support (PoE++ and Hi-PoE)	7.7.0	7.7.0	See the following improvements related to support for IEEE 802.3bt: • PoE++ and Hi-PoE—Up to 90W per port. • Single- and dual-signature powered devices (PDs). • Power budgeting is done on a first-come, first-served basis. • Power budget fields were added to **show power inline** . New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** > **PoE** New/Modified commands: **show power inline**
For models with built-in-switches, subinterfaces can't use VLAN 1	7.6	7.6	For models with built-in switches, you cannot create a subinterface using VLAN 1. VLAN 1 is reserved for the logical VLAN interface for switch ports. If you upgrade a 1010 to 7.6 or later, and you have assigned VLAN 1 to a subinterface, you must change the VLAN ID for your subinterface to a new VLAN. You will not be able to deploy the configuration using VLAN 1.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Secure Firewall supports Dual-Arm Deployment Mode on AWS with GWLB	7.6	7.6	The Secure Firewall supports the dual-arm deployment mode on AWS with GWLB. This mode enables the firewall to directly forward internet-bound traffic to the internet through the internet gateway after traffic inspection, while also performing network address translation (NAT).
Secure Firewall 1210/1220 hardware switch support	7.6	7.6	The Secure Firewall 1210/1220 supports setting each Ethernet interface to be a switch port or a firewall interface
Secure Firewall 1210CP PoE+ support on Ethernet ports 1/5-1/8	7.6	7.6	The Secure Firewall 1210CP supports Power over Ethernet+ (PoE+) on Ethernet ports 1/5-1/8.
VXLAN VTEP IPv6 support	7.4	Any	You can now specify an IPv6 address for the VXLAN VTEP interface. IPv6 is not supported for the threat defense virtual cluster control link or for Geneve encapsulation. New/Modified screens: • **Devices** > **Device Management** > **Edit** > **VTEP** > **Add VTEP** **Devices** > **Device Management** > **Edit** > **Interfaces** > **Add Interfaces** > **VNI Interface** Requires Firewall Threat Defense version 7.4. You can specify an IPv6 address for the VXLAN VTEP interface for Secure Firewall 4200 Series devices only.
Loopback interface support for BGP and management traffic	7.4	Any	You can use a loopback interface for: AAA, BGP, DNS, HTTP, ICMP, IPsec Flow Offload, NetFlow, SNMP, SSH, and Syslog. Requires Firewall Threat Defense version 7.4.
Loopback interface support for VTI	7.3	Any	You can now add a loopback interface. The loopback interface helps to overcome path failures. If an interface goes down, you can access all interfaces through the IP address assigned to the loopback interface. For VTI, in addition to setting a loopback interface as the source interface, support has also been added to inherit the IP address from a loopback interface instead of a statically configured IP address. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Add Interfaces** > **Add Loopback Interface**

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
IPv6 DHCP	7.3	Any	The Firewall Threat Defense now supports the following features for IPv6 addressing: DHCPv6 Address client, DHCPv6 Prefix Delegation client, BGP router advertisement for delegated prefixes, and DHCPv6 stateless server. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** > **Add/Edit Interfaces** > **IPv6** > **DHCP** • **Objects** > **Object Management** > **DHCP IPv6 Pool** New/Modified commands: **show bgp ipv6 unicast** , **show ipv6 dhcp** , **show ipv6 general-prefix**
Paired proxy VXLAN for the Firewall Threat Defense Virtual for the Azure Gateway Load Balancer	7.3	Any	You can configure a paired proxy mode VXLAN interface for the Firewall Threat Defense Virtual in Azure for use with the Azure Gateway Load Balancer (GWLB). New/Modified screens: **Devices** > **Device Management** > **Device** > **Interfaces** > **Add Interfaces** > **VNI Interface** Supported platforms: Firewall Threat Defense Virtual in Azure
VXLAN support	7.2	Any	VXLAN encapsulation support was added. New/Modified screens: **Devices** > **Device Management** > **Device** > **VTEP** , **Devices** > **Device Management** > **Device** > **Interfaces** > **Add Interfaces** > **VNI Interface** , and **Devices** > **Device Management** > **Device** > **Interfaces** edit physical interface **> General** Supported platforms: All.
Geneve support for the Firewall Threat Defense Virtual	7.1	Any	Geneve encapsulation support was added for the Firewall Threat Defense Virtual to support single-arm proxy for the Amazon Web Services (AWS) Gateway Load Balancer. This feature requires Snort 3.
31-bit Subnet Mask	7.0	Any	For routed interfaces, you can configure an IP address on a 31-bit subnet for point-to-point connections. This feature is not supported for BVIs for bridge groups or with multicast routing. New/Modified screens: **Devices** > **Device Management** > **Interfaces**

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Synchronization between the Firewall Threat Defense operational link state and the physical link state for the Firepower 4100/9300	6.7	Any	The Firepower 4100/9300 chassis can now synchronize the Firewall Threat Defense operational link state with the physical link state for data interfaces. **Note** This feature is not supported for clustering, container instances, or Firewall Threat Defense with a Radware vDP decorator. It is also not supported for ASA. New/Modified Firepower Chassis Manager screens: **Logical Devices** > **Enable Link State** New/Modified FXOS commands: **set link-state-sync enabled** , **show interface expand detail**
Firepower 1010 hardware switch support	6.5	Any	The Firepower 1010 supports setting each Ethernet interface to be a switch port or a firewall interface.
Firepower 1010 PoE+ support on Ethernet 1/7 and Ethernet 1/8	6.5	Any	The Firepower 1010 supports Power over Ethernet+ (PoE+) on Ethernet 1/7 and Ethernet 1/8 when they are configured as switch ports.
VLAN subinterfaces for use with container instances	6.3.0	Any	To provide flexible physical interface use, you can create VLAN subinterfaces in FXOS and also share interfaces between multiple instances.
Data-sharing interfaces for container instances	6.3.0	Any	To provide flexible physical interface use, you can share interfaces between multiple instances.
Integrated Routing and Bridging	6.2.0	Any	Integrated Routing and Bridging provides the ability to route between a bridge group and a routed interface. The bridge group participates in routing by using a Bridge Virtual Interface (BVI) to act as a gateway for the bridge group.



CHAPTER 43

Inline Sets and Passive Interfaces

You can configure IPS-only passive interfaces, passive ERSPAN interfaces, and inline sets.

- IPS interfaces, on page 1119
- Requirements and prerequisites for inline sets, on page 1123
- Guidelines for configuring inline sets and passive interfaces, on page 1125
- Configure a passive interface, on page 1127
- Configure Inline Sets, on page 1128

IPS interfaces

Intrusion Prevention System (IPS) interfaces include passive interfaces, passive Enhanced Remote Switch Port Analyzer (ERSPAN) interfaces, and inline sets. Interfaces in IPS-only mode bypass many firewall checks and support only IPS security policy (Snort). Implement IPS-only interfaces if you have a separate firewall protecting these interfaces and want to avoid firewall function overhead.

Note The firewall mode only affects regular firewall interfaces, and not IPS-only interfaces such as inline sets or passive interfaces. IPS-only interfaces can be used in both firewall modes.

Inline sets

An inline set acts like a bump on the wire and binds one or more interface pairs together to slot into an existing network. This function allows the Firewall Threat Defense to be installed in any network environment without requiring configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally and retransmit all traffic out of the other interface in the inline pair unless traffic is explicitly dropped. When you have multiple inline pairs in an inline set, traffic can only pass between the interfaces in the pair; it can't pass between interfaces in different pairs.

Tap mode functionality

With tap mode, the Firewall Threat Defense is deployed inline, but the network traffic flow is undisturbed. Instead, the Firewall Threat Defense makes a copy of each packet so that it can analyze the packets. Note that rules of these types do generate intrusion events when they are triggered, and the table view of intrusion events indicates that the triggering packets would have dropped in an inline deployment.

There are benefits to using tap mode with Firewall Threat Defense that are deployed inline:

- You can set up the cabling between the Firewall Threat Defense and the network as if the Firewall Threat Defense were inline and analyze the kinds of intrusion events the Firewall Threat Defense generates.
- Based on the results, you can modify your intrusion policy and add the drop rules that best protect your network without impacting its efficiency.
- When you are ready to deploy the Firewall Threat Defense inline, you can disable tap mode and begin dropping suspicious traffic without having to reconfigure the cabling between the Firewall Threat Defense and the network.



Note Tap mode *significantly* impacts Firewall Threat Defense performance, depending on the traffic.



Note Inline sets might be familiar to you as "transparent inline sets," but the inline interface type is unrelated to the transparent firewall mode or the firewall-type interfaces.

Multiple inline pairs and asynchronous routing

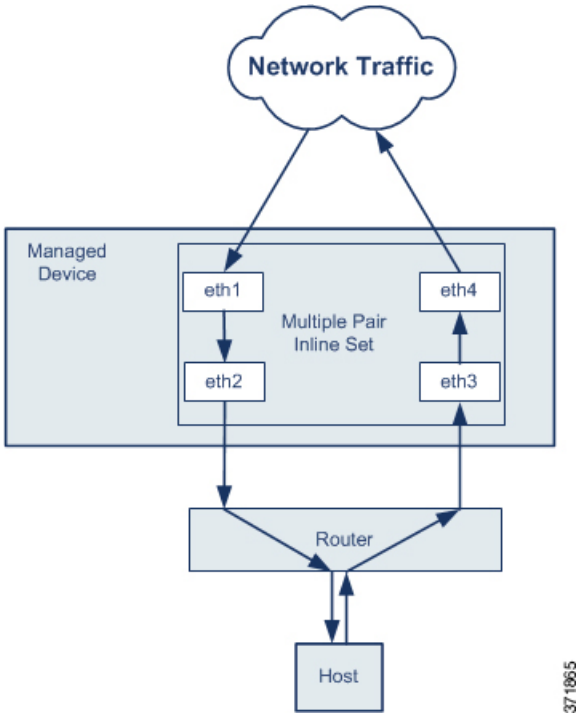
Asynchronous routing is a network configuration that routes traffic between a host on your network and external hosts through different inline pairs, depending on whether the traffic is inbound or outbound. It requires multiple inline pairs in an inline set to correctly analyze network traffic.

If you deploy asynchronous routing, but you include only one inline pair in an inline set, the device might not correctly analyze your network traffic because it might see only half of the traffic.

Multiple inline pairs configuration

Adding multiple inline pairs to the same inline set lets the system identify the inbound and outbound traffic as part of the same traffic flow. For passive interfaces only, you can also achieve this by including the interface pairs in the same security zone.

Figure 267: Asynchronous routing



Note If you assign multiple inline pairs to a single inline set, but you experience issues with duplicate traffic, you might need to reassign your inline pairs to separate inline sets or modify your security zones.

If fragments of a packet are received on different interface pairs, they are not reassembled and get dropped. Make sure all the fragments of the packet are received and sent on the same interface pair.

Passive interfaces

Passive interfaces monitor traffic flowing across a network using a switch SPAN or mirror port. The SPAN or mirror port allows for traffic to be copied from other ports on the switch. This function allows the system to view network traffic without participating in the flow of data.

Passive interface functionality

The SPAN or mirror port allows for traffic to be copied from other ports on the switch. When you configure the Firewall Threat Defense in a passive deployment, the Firewall Threat Defense cannot take certain actions such as blocking or shaping traffic.

Encapsulated remote switched port analyzer (ERSPAN) interfaces allow you to monitor traffic from source ports distributed over multiple switches and use GRE to encapsulate the traffic. ERSPAN interfaces are only allowed when the Firewall Threat Defense is in routed firewall mode.

Note Using SR-IOV interfaces as passive interfaces on NGFWv is not supported on some Intel network adapters (such as Intel X710 or 82599) using SR-IOV drivers due to a promiscuous mode restriction. In such cases, use a network adapter that supports this functionality. See Intel Ethernet Products for more information on Intel network adapters.

Hardware Bypass for inline sets

Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. It maintains network connectivity when software or hardware failures occur and is available for certain interface modules on supported models.

Supported interface modules

For information about which interface modules support this feature, see Requirements and prerequisites for inline sets, on page 1123.

Hardware Bypass triggers

Hardware Bypass can be triggered in these scenarios.

- Firewall Threat Defense crash
- Firewall Threat Defense reboot
- Security Module reboot
- Chassis crash
- Chassis reboot
- Manual trigger
- Chassis power loss
- Security Module power loss

Note Hardware bypass is intended for unplanned/unexpected failure scenarios, and is not automatically triggered during planned software upgrades. Hardware bypass activates only at the end of a planned upgrade process, when the Firewall Threat Defense application reboots.

Hardware bypass switchover

When you switch from normal operation to hardware bypass, or back to normal operation, you may experience several seconds of interrupted traffic. During this time, connections may drop.

Connection interruption factors

Several factors can affect the length of the traffic interruption during hardware bypass switchover.

- Copper port auto-negotiation
- Behavior of the optical link partner such as how it handles link faults and de-bounce timing
- Spanning tree protocol convergence
- Dynamic routing protocol convergence

You may also experience dropped connections due to application identification errors when analyzing connections midstream after the returning to normal operations.

Snort Fail Open vs. Hardware Bypass

Snort Fail Open and Hardware Bypass allow traffic to flow during system failures.

Functionality differences

You can configure the Snort Fail Open option for inline sets other than those in tap mode. This option either drops traffic or allows traffic to pass without inspection when the Snort process is busy or down. Snort Fail Open is supported on all inline sets except those in tap mode, not just on interfaces that support Hardware Bypass.

The Hardware Bypass functionality allows traffic to flow during a hardware failure, including a complete power outage, and certain limited software failures. A software failure that triggers Snort Fail Open does not trigger a Hardware Bypass.

Hardware Bypass status

When the system has power, the Bypass LED indicates the Hardware Bypass status. See the Firepower chassis hardware installation guide for LED descriptions.

Requirements and prerequisites for inline sets

This reference provides the requirements and prerequisites that must be met before configuring inline sets, including user role permissions and hardware bypass support information.

User roles

- Admin
- Access Admin
- Network Admin

Hardware Bypass Support

The Firewall Threat Defense supports Hardware Bypass for interface pairs on specific network modules on the following models:

- Firepower 2130 and 2140
- Secure Firewall 3100
- Firepower 4100

- Secure Firewall 4200
- Firepower 9300

Note The ISA 3000 has a separate implementation for Hardware Bypass, which you can enable using FlexConfig only (see FlexConfig policies, on page 1309). Do not use this chapter to configure ISA 3000 Hardware Bypass.

Note You can use Hardware Bypass interfaces as regular interfaces without the Hardware Bypass feature enabled.

The supported Hardware Bypass network modules for these models include:

- Firepower 2130 and 2140:
  - Firepower 6-port 1G SX FTW Network Module single-wide (FPR2K-NM-6X1SX-F)
  - Firepower 6-port 10G SR FTW Network Module single-wide (FPR2K-NM-6X10SR-F)
  - Firepower 6-port 10G LR FTW Network Module single-wide (FPR2K-NM-6X10LR-F)
- Secure Firewall 3100:
  - 6-port 1G SFP Fail-to-Wire Network Module, SX (multimode) (FPR3K-XNM-6X1SXF)
  - 6-port 10G SFP Fail-to-Wire Network Module, SR (multimode) (FPR3K-XNM-6X10SRF)
  - 6-port 10G SFP Fail-to-Wire Network Module, LR (single mode) (FPR3K-XNM-6X10LRF)
  - 6-port 25G SFP Fail-to-Wire Network Module, SR (multimode) (FPR3K-XNM-X25SRF)
  - 6-port 25G Fail-to-Wire Network Module, LR (single mode) (FPR3K-XNM-6X25LRF)
  - 8-port 1G Copper Fail-to-Wire Network Module, RJ45 (copper) (FPR3K-XNM-8X1GF)
- Secure Firewall 4200:
  - 6-port 1G SFP Fail-to-Wire Network Module, SX (multimode) (FPR4K-XNM-6X1SXF)
  - 6-port 10G SFP Fail-to-Wire Network Module, SR (multimode) (FPR4K-XNM-6X10SRF)
  - 6-port 10G SFP Fail-to-Wire Network Module, LR (single mode) (FPR4K-XNM-6X10LRF)
  - 6-port 25G SFP Fail-to-Wire Network Module, SR (multimode) (FPR4K-XNM-X25SRF)
  - 6-port 25G Fail-to-Wire Network Module, LR (single mode) (FPR4K-XNM-6X25LRF)
  - 8-port 1G Copper Fail-to-Wire Network Module, RJ45 (copper) (FPR4K-XNM-8X1GF)
- Firepower 4100:
  - Firepower 6-port 1G SX FTW Network Module single-wide (FPR4K-NM-6X1SX-F)
  - Firepower 6-port 10G SR FTW Network Module single-wide (FPR4K-NM-6X10SR-F)
  - Firepower 6-port 10G LR FTW Network Module single-wide (FPR4K-NM-6X10LR-F)

- Firepower 2-port 40G SR FTW Network Module single-wide (FPR4K-NM-2X40G-F)
- Firepower 8-port 1G Copper FTW Network Module single-wide (FPR-NM-8X1G-F)

- Firepower 9300:
  - Firepower 6-port 10G SR FTW Network Module single-wide (FPR9K-NM-6X10SR-F)
  - Firepower 6-port 10G LR FTW Network Module single-wide (FPR9K-NM-6X10LR-F)
  - Firepower 2-port 40G SR FTW Network Module single-wide (FPR9K-NM-2X40G-F)

Hardware Bypass can only use the following port pairs:

- 1 & 2
- 3 & 4
- 5 & 6
- 7 & 8

Guidelines for configuring inline sets and passive interfaces

Firewall mode restrictions

Use ERSPAN interfaces only when the device is in routed firewall mode.

Clustering restrictions

Consider clustering limitations when configuring inline sets and passive interfaces.

- Link State Propagation for an inline set is not supported with clustering.

Multi-Instance mode restrictions

Use unshared physical interfaces or EtherChannels when configuring multi-instance mode.

- Multi-instance shared interfaces are not supported. You must use an unshared interface.
- Multi-instance chassis-defined subinterfaces are not supported. You must use a physical interface or EtherChannel.

General configuration guidelines

Consider these interface and traffic inspection requirements when configuring inline sets and passive interfaces.

- Inline sets and passive interfaces support only physical interfaces and EtherChannels. You cannot use VLANs or other virtual interfaces, including multi-instance chassis-defined subinterfaces.
- Because IPS interfaces do not support regular firewall protections, traffic must pass through the same Firewall Threat Defense to ensure all traffic is inspected according to the IPS security policy (Snort).
- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the Firewall Threat Defense when using inline sets. If there are two neighbors on either side of the Firewall Threat Defense

running BFD, then the Firewall Threat Defense will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

- For inline sets and passive interfaces, the Firewall Threat Defense supports up to two 802.1Q headers in a packet (also known as Q-in-Q support), with the exception of the Firepower 4100/9300, which only supports one 802.1Q header. **Note:** Firewall-type interfaces do not support Q-in-Q, and only support one 802.1Q header.
- Firewall Threat Defense does not support frames with Ethertype 0x8988. If a source in the Application Centric Infrastructure (ACI) fabric has Precision Time Protocol (PTP) enabled, ACI inserts Ethertype 0x8988 into data-plane frames before mirroring. Firewall Threat Defense drops frames with Ethertype 0x8988 and does not inspect them.

Hardware Bypass Guidelines

Consider Hardware Bypass limitations and clustering requirements when configuring inline sets.

- Hardware Bypass ports are supported only for inline sets.
- Hardware Bypass ports cannot be part of an EtherChannel.
- Hardware Bypass is not supported in high availability mode.
- Hardware Bypass ports are supported with intra-chassis clustering on the Firepower 9300. Ports are placed in Hardware Bypass mode when the last unit in the chassis fails. Inter-chassis clustering is not supported, because inter-chassis clustering only supports Spanned EtherChannels; Hardware Bypass ports cannot be part of an EtherChannel.
- If all modules in an intra-chassis cluster on the Firepower 9300 fail, then Hardware Bypass is triggered on the final unit, and traffic continues to pass. When units come back up, Hardware Bypass returns to standby mode. If rules matching application traffic are used, connections might be dropped and require reestablishment. Connections are dropped because state information is not retained on the cluster unit, and the unit cannot identify the traffic as belonging to an allowed application. To prevent a traffic drop, use a port-based rule instead of an application-based rule, if appropriate for your deployment.
- You can use Hardware Bypass interfaces as regular interfaces without the Hardware Bypass feature enabled.

Unsupported firewall features on IPS interfaces

These firewall features are not supported on IPS interfaces.

- DHCP server
- DHCP relay
- DHCP client
- TCP Intercept
- Routing
- NAT
- VPN
- Application inspection

- QoS
- NetFlow
- VXLAN

Configure a passive interface

Configure a passive or ERSPAN interface to enable network monitoring and traffic analysis capabilities on your security device.

This section describes how to:

- Enable the interface. By default, interfaces are disabled.
- Set the interface mode to Passive or ERSPAN. For ERSPAN interfaces, you will set the ERSPAN parameters and the IP address.
- Change the MTU. By default, the MTU is set to 1500 bytes. For more information about the MTU, see MTU, on page 1104.
- Set a specific speed and duplex (if available). By default, speed and duplex are set to Auto.

Note For the Secure Firewall Threat Defense on the FXOS chassis, you configure basic interface settings on the Firepower 4100/9300. See Configure a physical interface for more information.

Before you begin

- If you are using EtherChannels, add them according to Configure an EtherChannel, on page 1011.

Follow these steps to configure a passive interface:

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 In the **Mode** drop-down list, choose **Passive** or **ERSPAN**.

Step 4 Enable the interface by checking the **Enabled** check box.

Set these values:

a. In the **Name** field, enter a name up to 48 characters in length.

b. From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

c. Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 5 (Optional) On **General**, set the **MTU** between 64 and 9198 bytes; for the Secure Firewall Threat Defense Virtual and Secure Firewall Threat Defense on the FXOS chassis, the maximum is 9000 bytes.

The default is 1500 bytes.

Step 6 For ERSPAN interfaces, set the following parameters:

- **Flow ID**—Configure the ID used by the source and destination sessions to identify the ERSPAN traffic, between 1 and 1023. This ID must also be entered in the ERSPAN destination session configuration.
- **Source IP**—Configure the IP address used as the source of the ERSPAN traffic.
- **IPv4**—Set the IPv4 address and mask

Step 7 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

The exact speed and duplex options depend on your hardware.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default.
- **Speed**—Choose **10**, **100**, **1000**, or **Auto**. Auto is the default.

Step 8 Click **OK**.

Step 9 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Inline Sets

Enable traffic inspection by configuring multiple inline sets using physical interfaces or EtherChannels, including optional Hardware Bypass.

Inline sets allow you to designate pairs of interfaces or EtherChannels for traffic inspection, with optional Hardware Bypass. You can configure multiple inline pairs per set.

Note For the Firepower 4100/9300 , you configure basic interface settings in FXOS on the chassis. See Configure a physical interface for more information.

Before you begin

- If you are using EtherChannels, add them according to Configure an EtherChannel, on page 1011 .
- We recommend that you set STP PortFast for STP-enabled switches that connect to the Firewall Threat Defense inline pair interfaces. This setting is especially useful for Hardware Bypass configurations and can reduce bypass times.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 For each interface in the inline pair, name the interface and enable it. Also configure necessary hardware settings, ensuring they match between interfaces in the inline pair. You can configure multiple pairs of interfaces.

a) Click **Edit** (✎) for the interface you want to edit.
b) In the **Name** field, enter a name up to 48 characters in length.
c) Enable the interface by checking the **Enabled** check box.
d) (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

e) Leave the **Mode** as **None** .

After you add this interface to an inline set, this field will show Inline for the mode.

f) Do not set the security zone yet; you must set it after you create the inline set later in this procedure.
g) (Optional) Set the duplex and speed by clicking **Hardware Configuration** .

The exact speed and duplex options depend on your hardware.

- **Duplex** —Choose **Full** , **Half** , or **Auto** . Auto is the default.
- **Speed** —Choose **10** , **100** , **1000** , or **Auto** . Auto is the default.

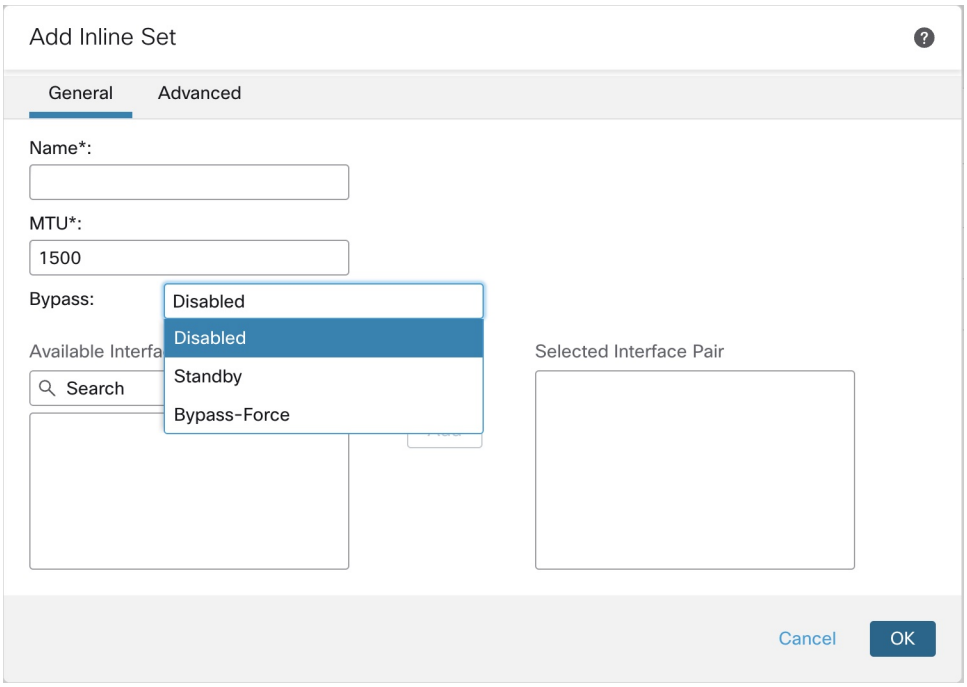
h) Click **OK** .

Do not set any other settings for this interface.

Step 3 Click **Inline Sets**.

Step 4 Click **Add Inline Set**.

Figure 268: Add Inline Set



The **Add Inline Set** dialog box appears with **General** selected.

Step 5 In the **Name** field, enter a name for the set.

Step 6 (Optional) Change the **MTU** to enable jumbo frames.

For inline sets, the MTU setting is not used. However, the jumbo frame setting is relevant to inline sets; jumbo frames enable the inline interfaces to receive packets up to 9000 bytes. To enable jumbo frames, you must set the MTU of *any* interface on the device above 1500 bytes.

Step 7 Configure Hardware Bypass.

a) For the **Bypass** mode, choose one of the following options:

- **Disabled** —Set Hardware Bypass to disabled for interfaces where Hardware Bypass is supported, or use interfaces where Hardware Bypass is not supported.
- **Standby** —Set Hardware Bypass to the standby state on supported interfaces. Only pairs of Hardware Bypass interfaces are shown. In the standby state, the interfaces remain in normal operation until there is a trigger event.
- **Bypass-Force** —Manually forces the interface pair to go into a bypass state. **Inline Sets** shows **Yes** for any interface pairs that are in Bypass-Force mode.

b) In the **Available Interfaces Pairs** area, click a pair and then click **Add** to move it to the **Selected Interface Pair** area.

All possible pairings between named and enabled interfaces with the mode set to None show in this area.

Step 8 (Optional) Click **Advanced** to set the following optional parameters:

- **Tap Mode** —Set to inline tap mode.

Note that you cannot enable this option and strict TCP enforcement on the same inline set.

Note
Perform Tap mode changes during a maintenance window. Changing the mode while the device is passing traffic can cause traffic disruption.

Note
Tap mode significantly impacts the Firewall Threat Defense performance, depending on the traffic.

- **Propagate Link State** —Configure link state propagation.

  Link state propagation automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up, also. In other words, if the link state of one interface changes, the device senses the change and updates the link state of the other interface to match it. Note that devices require up to 4 seconds to propagate link state changes. Link state propagation is especially useful in resilient network environments where routers are configured to reroute traffic automatically around network devices that are in a failure state.

  Note
  Do not enable **Propagate Link State** when using clustering.

- **Snort Fail Open** —Enable or disable either or both of the **Busy** and **Down** options if you want new and existing traffic to pass without inspection (enabled) or drop (disabled) when the Snort process is busy or down.

  By default, traffic passes without inspection when the Snort process is down, and drops when it is busy.

  When the Snort process is:

  - Busy—It cannot process traffic fast enough because traffic buffers are full, indicating that there is more traffic than the device can handle, or because of other software resource issues.
  - Down—It is restarting because you deployed a configuration that requires it to restart. See Configurations that restart the snort process when deployed or activated, on page 198 .

  When the Snort process is down and comes back up, it inspects *new* connections. To prevent false positives and false negatives, it does not inspect existing connections on inline, routed, or transparent interfaces because initial session information might have been lost while it was down.

  Note
  When Snort fails open, features that rely on the Snort process do not function. These include application control and deep inspection. The system performs only basic access control using simple, easily determined transport and network layer characteristics.

  Note
  The **Strict TCP Enforcement** option is not supported.

Step 9 Set the security zone for each interface.

a) Click **Interfaces** .

b) Click **Edit** ( ) for the member interface.

c) From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New** .

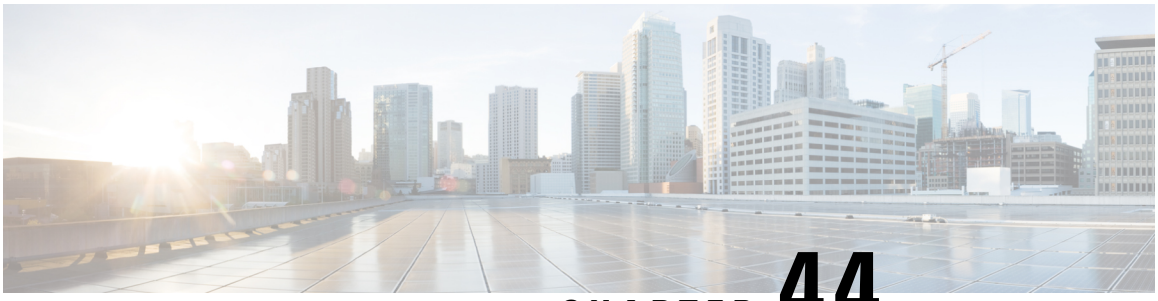
You can only set the zone after you add the interface to the inline set; adding it to an inline set configures the mode to Inline and lets you choose inline-type security zones.

d) Click **OK** .

Click **Save**.

What to do next

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.



CHAPTER 44

Device Settings

After you add a device, you can edit device-related settings on the **Device** page.

1. Choose **Devices** > **Device Management**.
2. Next to the device you want to modify, click **Edit** ( ).
3. Click **Device**.

- Edit general settings, on page 1133
- Edit license settings, on page 1140
- View system information, on page 1141
- View the inspection engine, on page 1142
- Edit health settings, on page 1143
- Edit management settings, on page 1153
- View inventory details, on page 1186
- Edit applied policies, on page 1186
- Edit advanced settings, on page 1188
- Edit deployment settings, on page 1192
- Edit cluster health monitor settings, on page 1195
- Hot swap an SSD, on page 1198
- Disable the USB port, on page 1201
- Configure SNMP for FXOS, on page 1204
- Configure alarms for the ISA 3000, on page 1207
- History for Device Settings, on page 1219

Edit general settings

Edit general device settings to customize the device name, packet transfer options, and force deployment settings for managed devices.

The **General** section of the **Device** page displays the settings described in this table.

Figure 269: General

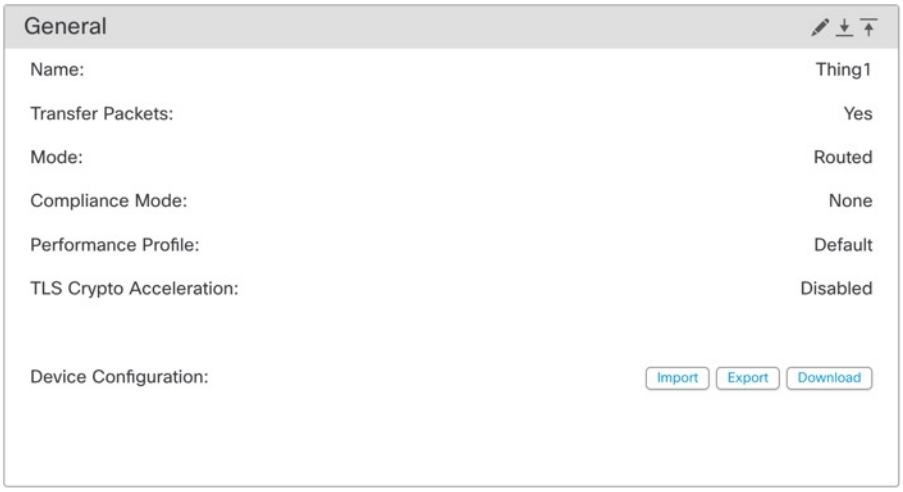


Table 70: General section table fields

Field	Description
Name	The display name of the device on the Cloud-Delivered Firewall Management Center.
Transfer Packets	This displays whether or not the managed device sends packet data with the events to the Cloud-Delivered Firewall Management Center.
Mode	The displays the mode of the management interface for the device: **routed** or **transparent**.
Compliance Mode	This displays the security certifications compliance for a device. Valid values are CC, UCAPL and None.
Performance Profile	This displays the core allocation performance profile for the device, as configured in the platform settings policy.
TLS Crypto Acceleration:	Shows whether TLS crypto acceleration is enabled or disabled.
Device Configuration	Lets you copy, export, or import a configuration. See Copy a configuration to another device, on page 1135 and Export and import the device configuration, on page 1136.
OnBoarding Method	Shows whether the device was registered using a registration key or using the serial number (zero-touch provisioning).

You can edit some of these settings from this section.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **General** section, click **Edit** ( ).

a) Enter a **Name** for the managed device.

b) Check **Transfer Packets** to allow packet data to be stored with events on the Cloud-Delivered Firewall Management Center.

c) Click **Force Deploy** to force deployment of current policies and device configuration to the device.

Note
Force-deploy consumes more time than the regular deployment since it involves the complete generation of the policy rules to be deployed on the Firewall Threat Defense.

Step 5 For **Device Configuration** actions, see Copy a configuration to another device, on page 1135 and Export and import the device configuration, on page 1136.

Step 6 Click **Deploy**.

What to do next

- Deploy configuration changes.

Copy a configuration to another device

Copy configurations and policies from a pre-configured device to a new device, instead of manually reconfiguring the new device.

When a new device is deployed in the network, you can copy configurations and policies from a pre-configured device, instead of manually reconfiguring the new device.

Before you begin

Confirm that:

- The source and destination devices are the same model and are running the same version of the software.
- The source is either a standalone device or a high availability pair.
- The destination device is a standalone device.
- The source and destination devices have the same number of physical interfaces.
- The source and destination devices are in the same firewall mode: routed or transparent.
- The source and destination devices are in the same security-certifications-compliance mode.
- The source and destination devices are in the same domain.
- No configuration deployment is in progress on the source or the destination devices.

Follow these steps to copy a configuration to another device:

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **General** section, do one of the following:

- Click **Get Device Configuration** ( ) to copy device configuration from another device to the new device. On the **Get Device Configuration** page, select the source device in the **Select Device** drop-down list.
- Click **Push Device Configuration** ( ) to copy device configuration from the current device to the new device. On the **Push Device Configuration** page, select the destination to which configuration is to be copied in the **Target Device** drop-down list.

Step 5 (Optional) Check **Include shared policies configuration** check box to copy policies.

Shared policies such as AC policy, NAT, Platform Settings and FlexConfig policies can be shared across multiple devices.

Step 6 Click **OK**.

You can monitor the status of the copy device configuration task on **Tasks** in the Message Center.

When the copy device configuration task is initiated, it erases the configuration on the target device and copies the configuration of the source device to the destination device.

Warning When you have completed the copy device configuration task, you cannot revert the target device to its original configuration.

Export and import the device configuration

Export device-specific configuration to create a backup that can be imported for the same device in various scenarios including moving the device to a different Cloud-Delivered Firewall Management Center, restoring old configurations, or reregistering a device.

You can export all device-specific configurations available on the Device pages, including:

- Interfaces
- Inline Sets
- Routing
- DHCP
- VTEP

- Associated objects

You can then import the saved configuration for the same device in these use cases:

- Moving the device to a different Cloud-Delivered Firewall Management Center—First unregister the device from the original Cloud-Delivered Firewall Management Center, then add the device to the new Cloud-Delivered Firewall Management Center. Then you can import the saved configuration.
- Restore an old configuration—If you deployed changes that negatively impacted the operation of the device, you can import a backup copy of a known working configuration to restore a previous operational state.
- Reregistering a device—If you unregister a device from the Cloud-Delivered Firewall Management Center, but then want to add it back, you can import the saved configuration.

Before you begin

Review these guidelines:

- You can only import the configuration to the same device (the UUID must match). You cannot import a configuration to a different device, even if it is the same model.
- Do not change the version running on the device between exporting and importing; the version must match.
- If you make inventory changes after your export (such as adding or deleting network modules or configuring or joining breakout ports), the device inventory will not match the Cloud-Delivered Firewall Management Center. In this case, the device inventory will be maintained, and you will be prompted to sync the interfaces (see Sync Interface Changes with the Cloud-Delivered Firewall Management Center, on page 1014) and discard incompatible configuration in the Cloud-Delivered Firewall Management Center when you try to deploy. You will have to repeat the inventory changes and related configuration in the Cloud-Delivered Firewall Management Center.
- If you export a standalone configuration, you cannot import it to a high availability pair or vice versa.
- If an object doesn't exist, it will be created. If an object exists, but the value is different, refer to this table:

Table 71: Object import action

Scenario	Import Action
Object exists with the same name and value.	Reuse existing objects.
Object exists with the same name but different value.	Network and Port objects: Create object overrides for this device. See Object overrides, on page 2939. Interface objects: Create new objects. For example, if both the type (security zone or interface group) and the interface type (routed or switched, for example) do not match, then a new object is created. All other objects: Reuse existing objects even though the values are different.
Object doesn't exist.	Create new objects.

Follow these steps to export and import the device configuration:

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to edit, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 Export the configuration.

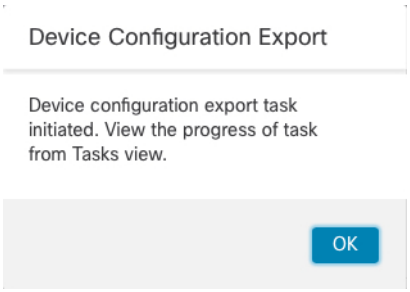
a) In the **General** area, click **Export**.

Figure 270: Export device configuration



You are prompted to acknowledge the export; click **OK**.

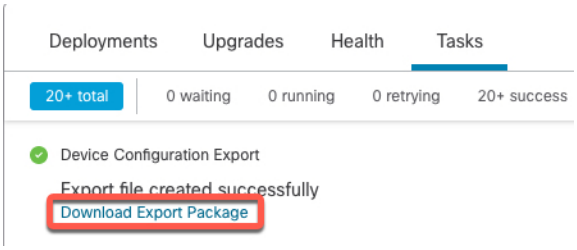
Figure 271: Acknowledge export



You can view the export progress in the **Tasks** page.

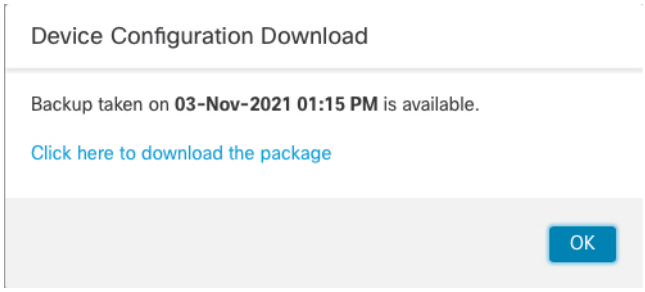
b) Click **Notifications**, and then click the **Tasks** tab. Verify if the export has completed, and then click **Download Export Package**. Alternatively, you can click the **Download** button in the **General** area.

Figure 272: Export task



You are prompted to download the package; click **Click here to download the package** to save the file locally, and then click **OK** to exit the dialog box.

Figure 273: Download package



Step 5 Import the configuration.

a) In the **General** area, click **Import**.

Figure 274: Import device configuration



You are prompted to acknowledge that the current configuration will be replaced. Click **Yes** and then navigate to the configuration package (with the suffix .sfo; note that this file is different from the Backup/Restore files).

Figure 275: Import package

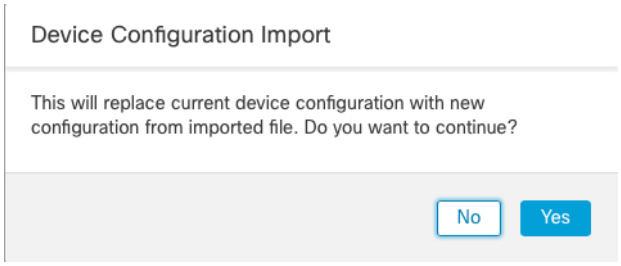
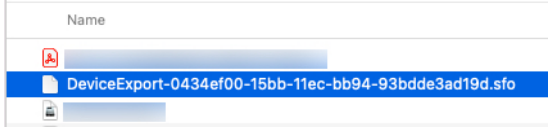
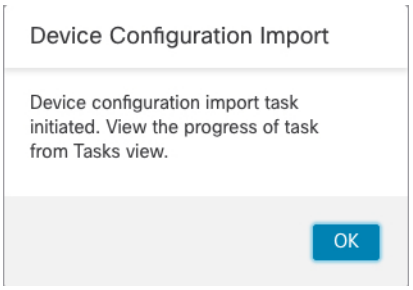


Figure 276: Navigate to package



You are prompted to acknowledge the import; click **OK**.

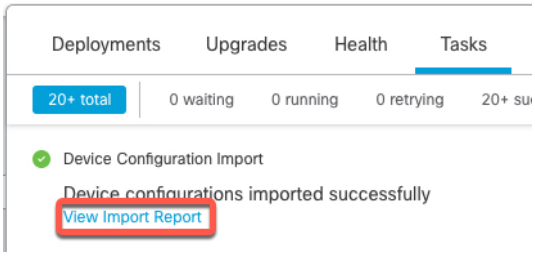
Figure 277: Acknowledge import



You can view the import progress in the **Tasks** page.

b) To view the import reports, click **Notifications**, and then click the **Tasks** tab. Click **View Import Report**.

Figure 278: View import report



The **Device Configuration Import Reports** page provides links to available reports.

Cisco Firepower Management Center

Device Configuration Import Reports

Device	Shared Policies	Device Configurations
0434ef00-15bb-11ec-bb94-93bdde3ad19d	Report does not exist	Device configurations import report

c) Deploy configuration changes.

Edit license settings

Enable or disable specific licenses on your managed device to control which features are available for use.

The **License** section of the **Device** page displays the licenses enabled for the device.

You can enable licenses on your device if you have available licenses on your Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to enable or disable licenses, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **License** section, click **Edit** ( ).

Step 5 Check or clear the check box next to the license you want to enable or disable for the managed device.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes.

View system information

The **System** section of the **Device** page displays a read-only table containing system information. You can also shut down or restart the device from this pane, using the icons at the top-right corner.

Figure 279: System

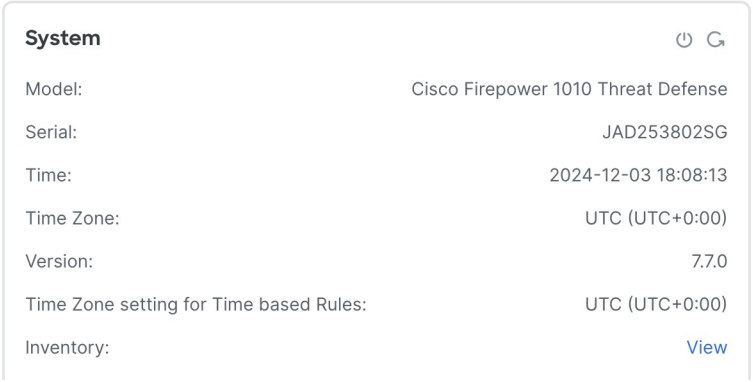


Table 72: System section table fields

Field	Description
Shut Down Device ( )	Shuts down the device. See Shut down or restart the device, on page 107.
Restart Device ( )	Restarts the device. See Shut down or restart the device, on page 107.
Model	Model name and number.

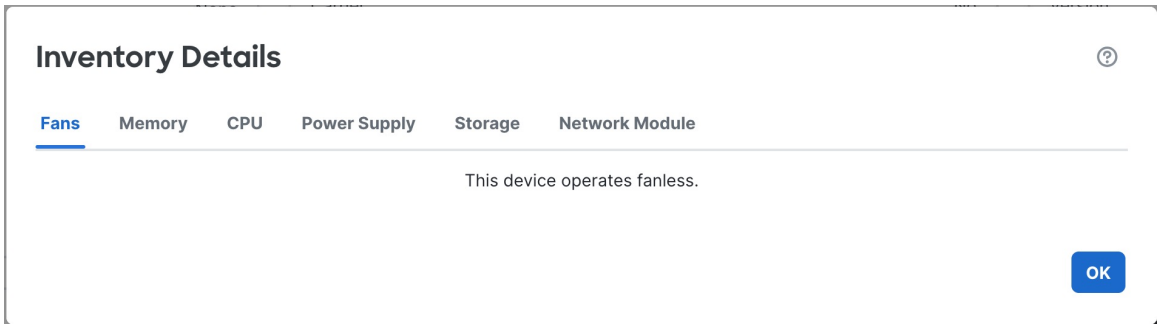
Field	Description
Serial	PCB (circuit board) serial number. The firewall includes two serial numbers: the chassis serial number and the PCB serial number. The chassis serial number is shown in the Inventory section.
Time	Current system time of the device.
Time Zone	Time zone.
Version	Version of the software currently installed on the managed device.
Time Zone setting for Time based rules	Current system time of the device in the time zone specified in device platform settings.
Inventory	Inventory details. See View Device Inventory.

View device inventory

Click **View** next to **Inventory** in the **System** section to view a table of device inventory information including Fans, Memory, CPU, Power Supply, Storage, and Network Modules.

The **Inventory Details** table displays information about all the Cisco products installed in the Firewall Threat Defense devices assigned with a product identifier (PID). The PID is the product name using which the product can be ordered.

Figure 280: Inventory details



View the inspection engine

The Inspection Engine section of the **Device** page shows whether your device uses Snort 2 or Snort 3. To switch the inspection engine, see *Enable Snort 3 on an Individual Device* in the *Custom Snort 3 Intrusion Policies for Access Control* .

Note Snort3 is default for Secure Firewall 1200. It does not support Snort2.

Edit health settings

The **Health** section of the **Device** page displays the information described in this table.

Figure 281: Health

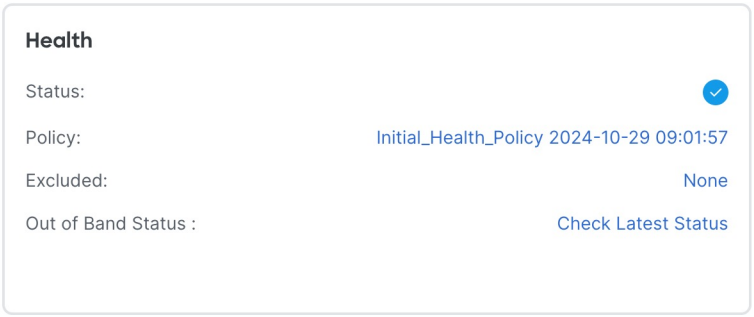


Table 73: Health section table fields

Field	Description
Status	An icon that represents the current health status of the device. Clicking the icon displays the Health Monitor for the appliance.
Policy	A link to a read-only version of the health policy currently deployed at the device.
Excluded	A link to the **Health Exclude** page, where you can enable and disable health exclusion modules.
Out of Band Status	A link to the **Out-of-Band configuration details** dialog box where you can view out-of-band configuration changes made at the device CLI. You must acknowledge the configuration differences and manually match any changes you want to keep in the Cloud-Delivered Firewall Management Center before the next deployment. See Out-of-band configuration detection, on page 1143.

Out-of-band configuration detection

Out-of-band configuration detection

- enables select configuration changes directly at the device CLI when you lose the management connection to your device
- allows restoration of the management connection if you are using a data interface for manager access
- detects configuration changes on the device after the management connection is restored, requiring manual acknowledgment and synchronization with the Cloud-Delivered Firewall Management Center.

Configuration change purposes

If you lose the management connection to your device, you can make select configuration changes directly at the device CLI to:

- Restore the management connection if you are using a data interface for manager access
- Make select configuration changes that can't wait until the connection is restored

Caution You are expected to know the commands that are required for recovery or emergency use. Do not use this feature to experiment with configuration changes. If you do not know which commands are required or are unsure about the effect of a command, we recommend that you contact Cisco TAC for guidance.

After the management connection is restored, the Cloud-Delivered Firewall Management Center will detect the configuration changes on the device. It does not automatically update the device configuration in the Cloud-Delivered Firewall Management Center.You must view the configuration differences, acknowledge that the device configuration differs, and manually apply the same changes in the Cloud-Delivered Firewall Management Center before deploying.

Caution When you deploy after acknowledgment, any configurations not present in the Cloud-Delivered Firewall Management Center configuration will be overwritten on the device.

Guidelines for out-of-band configuration

Supported feature areas in recovery-config mode

You can configure the following feature areas at the diagnostic CLI in recovery-config mode:

- Interfaces
- Static Routes
- Dynamic Routing: BGP and OSPF
- Prefilters
- Site-to-site VPN

Like other diagnostic CLI commands, refer to the ASA command reference for more information about each command.

Unsupported features

Be aware of these feature limitations in recovery-config mode:

- Not supported in multi-instance mode.
- You cannot add or delete EtherChannels.
- Some platform-dependent interface commands such as speed, duplex, and fec are not supported.

High availability and clustering restrictions

Consider these restrictions when using recovery-config mode in high availability or clustering environments:

- Recovery-config mode is only available on the active/control node.
- If a failover or cluster switchover occurs before you exit the recovery-config-mode session, the Cloud-Delivered Firewall Management Center will not detect the change on the new active/control node. We recommend re-entering recovery-config mode on the new active/control node and making a small change to trigger discovery of all of your previous changes. Otherwise, if you do not manually match the changes in the Cloud-Delivered Firewall Management Center, they will be overwritten at deployment without any notification.
- If you make out-of-band configuration changes on the active/control node and the high availability or cluster enters 'split brain' mode before a configuration sync (for example, due to a failover or cluster-control-link failure), the changes will be lost if a different node becomes active/control after the cluster returns to a healthy state.
- If you have an active recovery-config-mode session, then new nodes cannot join or rejoin the high availability/cluster until you exit the session.

Additional configuration guidelines

Follow these guidelines when making configuration changes in recovery-config mode:

- To modify an existing rule or route, delete the existing command using the **no** form of the command and then re-add the modified rule. This method avoids conflicts and errors. For example:

  Incorrect:

```
firepower# show running-config route
route outside 10.0.0.0 255.0.0.0 20.1.1.1 1
firepower# configure recovery-config

 CAUTION: The config CLI is for emergency use only. Use the config CLI if the management
 center is
unreachable, and use it only under exceptional circumstances, such as loss of connectivity
 or
to restore manager access. Do not change management center's auto-generated
configurations.

 After your management center is reachable, manually make the same configuration changes
 in the
management center. The management center cannot implement them automatically. When you
 deploy
from the management center, out-of-band configuration changes will be overwritten. Also,
 node join
will be blocked till config CLI session is active, so make sure to exit from the config
 CLI after
changes are made.

Would you like to proceed ? [Y]es/[N]o: y
firepower(recovery-config)# route outside 10.0.0.0 255.0.0.0 30.1.1.1
firepower(recovery-config)# exit
Unsaved changes are not kept if you reboot.Save changes to memory ? [Y]es/[N]o: y
Cryptochecksum: ccfc11a8 4e46d55e 0c99b5ae 3b18a8f1

3939 bytes copied in 0.70 secs
firepower# show running-config route
route outside 10.0.0.0 255.0.0.0 20.1.1.1 1
```

```
route outside 10.0.0.0 255.0.0.0 30.1.1.1 1
firepower#
```

In this case, a second route is added instead of replacing the first route.

Correct:

```
firepower# show running-config route
route outside 10.0.0.0 255.0.0.0 20.1.1.1 1
firepower# configure recovery-config

 CAUTION: The config CLI is for emergency use only. Use the config CLI if the management
 center is
unreachable, and use it only under exceptional circumstances, such as loss of connectivity
 or
to restore manager access. Do not change management center's auto-generated
configurations.

 After your management center is reachable, manually make the same configuration changes
 in the
management center. The management center cannot implement them automatically. When you
 deploy
from the management center, out-of-band configuration changes will be overwritten. Also,
 node join
will be blocked till config CLI session is active, so make sure to exit from the config
 CLI after
changes are made.

Would you like to proceed ? [Y]es/[N]o: y
firepower(recovery-config)# no route outside 10.0.0.0 255.0.0.0 20.1.1.1
firepower(recovery-config)# route outside 10.0.0.0 255.0.0.0 30.1.1.1
firepower(recovery-config)# exit
Unsaved changes are not kept if you reboot.Save changes to memory ? [Y]es/[N]o: y
Cryptochecksum: 81bcc51d 43771bbd 15b6dde6 afeb3442

3945 bytes copied in 0.70 secs
firepower# show running-config route
route outside 10.0.0.0 255.0.0.0 30.1.1.1 1
firepower#
```

- If you have auto rollback enabled (refer to Edit deployment settings, on page 1192), and you lose management connectivity because of a deployment, you should not start an out-of-band configuration. Instead, either wait 20 minutes for auto rollback to the previous deployment to occur or manually roll back at the CLI using the **configure policy rollback** command (refer to Manually Roll Back the configuration if the Cloud-Delivered Firewall Management Center loses Connectivity, on page 1181). Auto rollback will overwrite out-of-band configuration changes if the management connection is still down.
- For prefilter rules, we don't recommend adding completely new rules (the **access-control advanced** command). Integration of prefilter rules with the intrusion policy and logging requires the Cloud-Delivered Firewall Management Center, which generates the rule ID and integrates it with other policies.
- All recovery-config-mode sessions will be logged in syslog with the username "enable_15".

Access recovery-config mode in the diagnostic CLI

This task allows you to make out-of-band configuration changes when the management connection is down using the diagnostic CLI recovery-config mode.

You can use the diagnostic CLI recovery-config mode to make out-of-band configuration changes when the management connection is down. Be sure to make the same changes in the Cloud-Delivered Firewall Management Center. Local changes will always be overwritten by the Cloud-Delivered Firewall Management Center deployment.

For high availability and clustering, make your changes on the active/control node. This mode is not supported in multi-instance mode.

Procedure

Step 1 Connect to the device CLI using either the console port or SSH.

See Log into the Command-Line Interface on the device, on page 99.

Step 2 Access the diagnostic CLI.

system support diagnostic-CLI

enable (Press enter without entering a password when prompted.)

Example:

```
> system support diagnostic-cli
firepower> enable
Password:
```

Step 3 Show the current running configuration for reference.

show running-config

Note
You cannot enter **show** commands in recovery-config mode.

Step 4 Enter recovery-config mode.

configure recovery-config

Example:

```
firepower# configure recovery-config

 CAUTION: The config CLI is for emergency use only. Use the config CLI if the management
center is
unreachable, and use it only under exceptional circumstances, such as loss of connectivity
 or
to restore manager access. Do not change management center's auto-generated configurations.


 After your management center is reachable, manually make the same configuration changes
in the
management center. The management center cannot implement them automatically. When you
deploy
from the management center, out-of-band configuration changes will be overwritten. Also,
node join
will be blocked till config CLI session is active, so make sure to exit from the config CLI
 after
changes are made.

Would you like to proceed ? [Y]es/[N]o: y
```

```
firepower(recovery-config)#
```

Step 5 You can now enter select configuration commands.

Enter **?** to view available commands.

See Guidelines for out-of-band configuration, on page 1144 for supported feature areas.

See the ASA configuration guides or command reference for details about the commands.

Tip
Keep track of all of the commands you changed. Although the Cloud-Delivered Firewall Management Center shows you the differential later, it's good practice to keep a record of your command changes in case you need to make iterative changes for restoring the management connection.

Example:

```
firepower(recovery-config)# ?

  access-list           Configure an access control element
  as-path               BGP autonomous system path filter
  bfd                   BFD configuration commands
  bfd-template          BFD template configuration
  cluster               Cluster configuration
  community-list        Add a community list entry
  crypto                Configure IPSec, ISAKMP, Certification authority, key
  end                   Exit from configure mode
  exit                  Exit from config mode
  extcommunity-list     Add a extended community list entry
  group-policy          Configure or remove a group policy
  interface             Select an interface to configure
  ip                    Configure IP address pools
  ipsec                 Configure transform-set, IPSec SA lifetime and PMTU
                        Aging reset timer
  ipv6                  Configure IPv6 address pools
  ipv6                  Global IPv6 configuration commands
  isakmp                Configure ISAKMP options
  jumbo-frame           Configure jumbo-frame support
  mac-address           MAC address options
  management-interface  Management interface
  mtu                   Specify MTU(Maximum Transmission Unit) for an interface

  no                    Negate a command or set its defaults

  policy-list           Define IP Policy list
  prefix-list           Build a prefix list
  route                 Configure a static route for an interface
  route-map             Create route-map or enter route-map configuration mode
  router                Enable a routing process
  sla                   IP Service Level Agreement
  sysopt                Set system functional options
  time-range            Define time range entries
  tunnel-group          Create and manage the database of connection specific
                        records for IPSec connections
  vpdn                  Configure VPDN feature
  vrf                   Configure a VRF
  zone                  Create or show a Zone
firepower(recovery-config)#
```

Step 6 Exit recovery-config mode to be prompted to save your changes. Enter **exit** to exit each submode until you return to enable mode.

You can choose to save your changes to the startup configuration or keep changes only in the running configuration by not saving. Running configuration changes won't be retained after a reboot. If you make additional changes later and decide to save the configuration, all of your previous changes are also saved, since the entire running configuration is saved.

Deployment is be blocked when the recovery-config-mode session is open.

Example:

```
firepower(recovery-config)# interface Ethernet0/1
firepower(config-if)# ip address 10.0.0.2 255.0.0.0
firepower(config-if)# exit
firepower(recovery-config)# exit
Unsaved changes are not kept if you reboot. Save changes to memory ? [Y]es/[N]o: y

Cryptochecksum: 81a9073e f9535916 9c333d7e 9a3e5e76

3756 bytes copied in 0.70 secs
firepower#

Unsaved changes are not kept if you reboot. Save changes to memory ? [Y]es/[N]o:

Cryptochecksum: 81a9073e f9535916 9c333d7e 9a3e5e76

3756 bytes copied in 0.70 secs
firepower#
```

Step 7 Return to the Firewall Threat Defense CLI by typing Ctrl+a, then d, or you can enter **exit** to exit each mode.

Note
If you type Ctrl+a, then d to return to the Firewall Threat Defense CLI without first exiting recovery-config mode, the recovery-config-mode session will remain open, and deployment will be blocked.

Example:

```
firepower# exit

Logoff

User enable_1 logged in to firepower
Logins over the last 1 days: 4.  Last login: 20:42:51 UTC Dec 4 2024 from console
Failed logins since the last login: 0.
Type help or '?' for a list of available commands.
firepower> exit
Console connection detached.
>
```

Acknowledge the out-of-band configuration

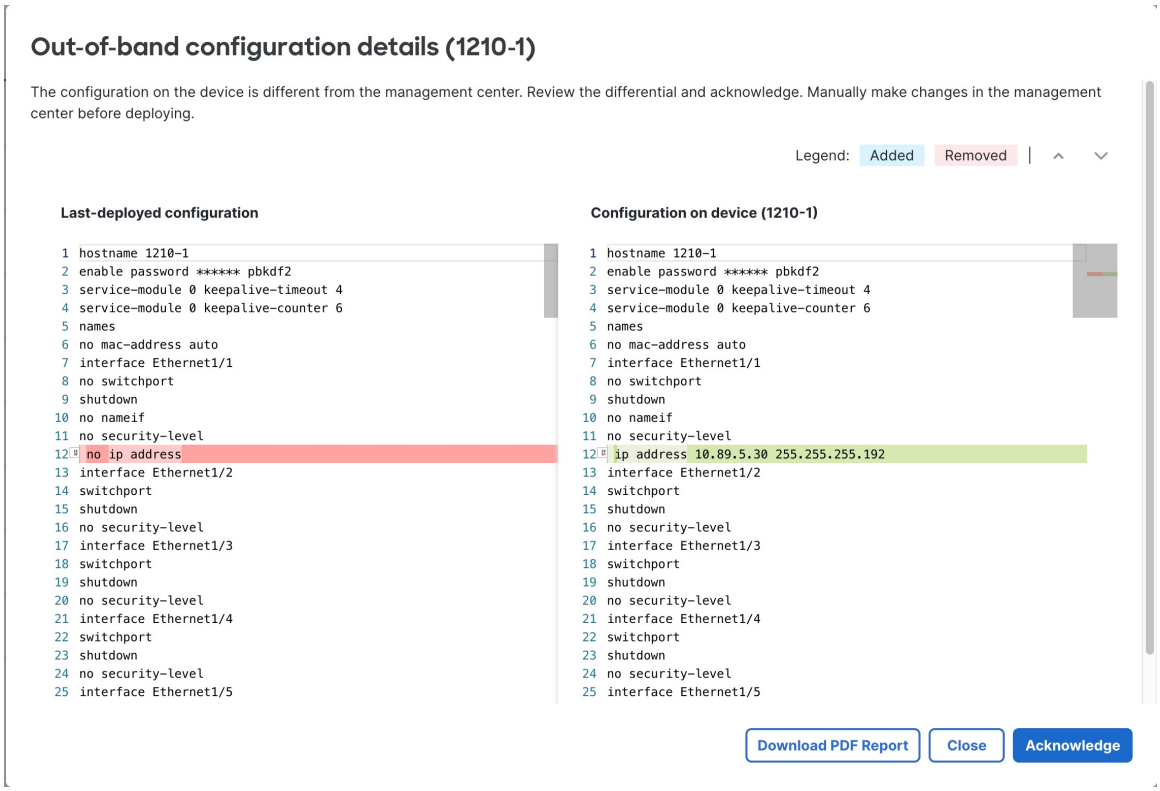
This task acknowledges out-of-band configuration changes detected by the Cloud-Delivered Firewall Management Center and enables deployment by matching configurations within the system.

When the Cloud-Delivered Firewall Management Center detects an out-of-band configuration change on a device, you must acknowledge the changes and match the configuration within the Cloud-Delivered Firewall Management Center that you want to keep. Until you acknowledge the changes, deployment will be blocked.

Procedure

Step 1 Open the **Out-of-Band configuration details** dialog box.

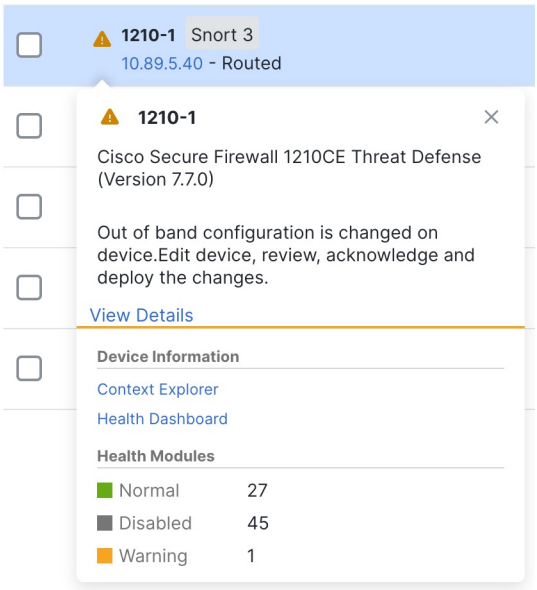
Figure 282: Out-of-Band configuration details



Note
Some commands, when set to a default setting, don't appear in the command output. However, the non-default command will show on either side as green (added) or red (removed). For example, if you add **no shutdown** to an interface in recovery-config mode, the **shutdown** command will show in red on the left **Last-deployed configuration** pane while **no shutdown** will *not* appear in the right **Configuration on device** pane. In this case, although the default setting for an interface is **shutdown**, the parser considers **no shutdown** to be the default and doesn't show it.

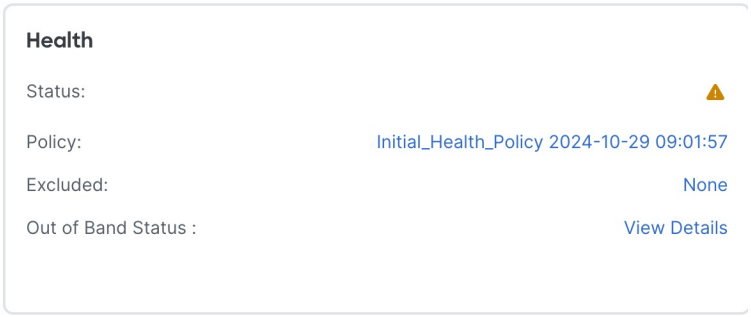
You can open the dialog box from multiple locations. For example, on the **Devices** > **Device Management** page, your device will have a warning. Click **View Details**.

Figure 283: Device management warning



Or, from the **Devices** > **Device Management**, and then navigate to **Health** tile under **Device** tab, you can click **View Details**.

Figure 284: Health Out-of-Band status



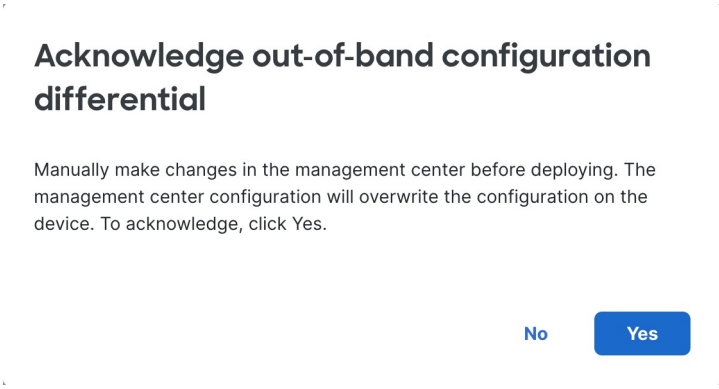
Note
If the out-of-band notification hasn't yet reached the Cloud-Delivered Firewall Management Center, you can check for changes using the **Check Latest Status** link in the Health tile.

Step 2 Click **Download PDF Report** so you can refer to the configuration changes you need to make after you close the dialog box.

Or you can bring up the dialog box at any time to review the changes.

Step 3 Click **Acknowledge**, and then **Yes**.

Figure 285: Acknowledge

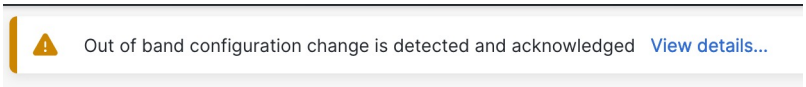


If you want to prevent an accidental deployment until after you've made your configuration changes, you can instead make the changes and then come back and click **Acknowledge**.

Step 4 Click **Close** on the **Out-of-Band configuration details** dialog box.

You can still revisit the dialog box to review the changes you need to make until you deploy. The status on the Device page changes to show you have acknowledged the out-of-band configuration:

Figure 286: Acknowledgement status

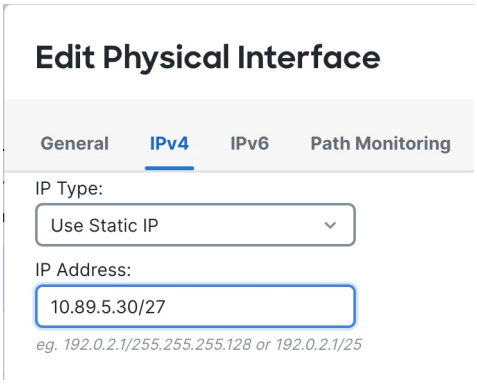


Step 5 Make the configuration changes that you made at the CLI.

You'll need to match the configuration CLI to Cloud-Delivered Firewall Management Center screens; there aren't links from the CLI changes directly to screens.

If you don't want to keep your changes, you can simply deploy and overwrite the device configuration. You should make all necessary changes to maintain the management connection as well as any other changes you want to keep. For example, if you changed the IP address at the CLI, you need to go to the **Interfaces** page, edit the interface, and set that IP address to match:

Figure 287: Match the IP address change



There is no checking mechanism that you made the same change; you could set the IP address differently if you want.

Step 6 Deploy configuration changes.

After you deploy, you can view the configuration differential—whether you made the changes or not—on the **System** (⚙) > **Monitoring** > **Audit** page. Check for the subsystem called *Device* > *Device Management* > *Out of band changes*.

Edit management settings

These settings control how the Cloud-Delivered Firewall Management Center establishes the management connection with the device.

Configure a redundant manager access data interface

This task configures a redundant manager access data interface to ensure continuous management connectivity. The secondary interface provides automatic failover capabilities when the primary management interface becomes unavailable, maintaining device management functionality.

When you use a data interface for manager access, you can configure a secondary data interface to take over management functions if the primary interface goes down. You can configure only one secondary interface. The device uses SLA monitoring to track the viability of the static routes and an ECMP zone that contains both interfaces so management traffic can use both interfaces.

Before you begin

- The secondary interface needs to be in a separate security zone from the primary interface.
- All of requirements apply to the primary interface also apply to the secondary interface.

Follow these steps to configure a redundant manager access data interface:

Procedure

Step 1 On the **Devices** > **Device Management** page, click **Edit** (✎) for the device.

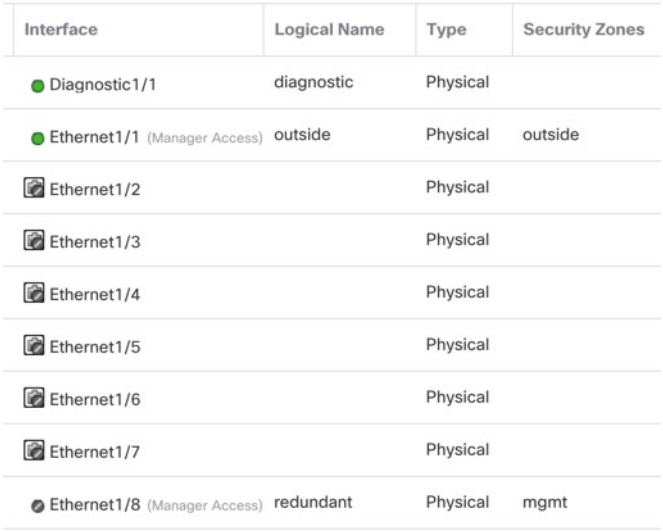
Step 2 Enable manager access for the secondary interface.

This setting is in addition to standard interface settings such as enabling the interface, setting the name, setting the security zone, and setting a static IPv4 address.

a) Choose **Interfaces** > **Edit Physical Interface** > **Manager Access**.
b) Check **Enable management on this interface for the Manager**.
c) Click **OK**.

Both interfaces show **(Manager Access)** in the interface listing.

Figure 288: Interface listing



Step 3 Add the secondary address to the **Management** settings.

a) Click **Device**, and view the **Management** area.

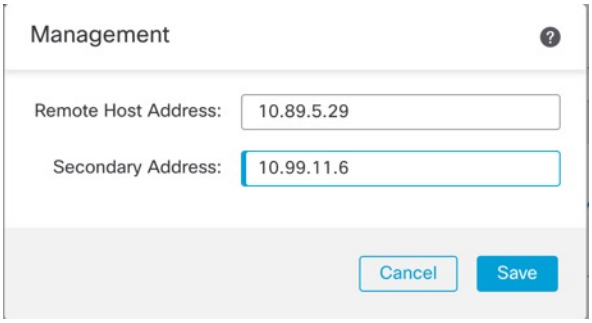
b) Click **Edit** ( ).

Figure 289: Edit management address



c) In the **Management** dialog box, modify the name or IP address in the **Secondary Address** field

Figure 290: Management IP address

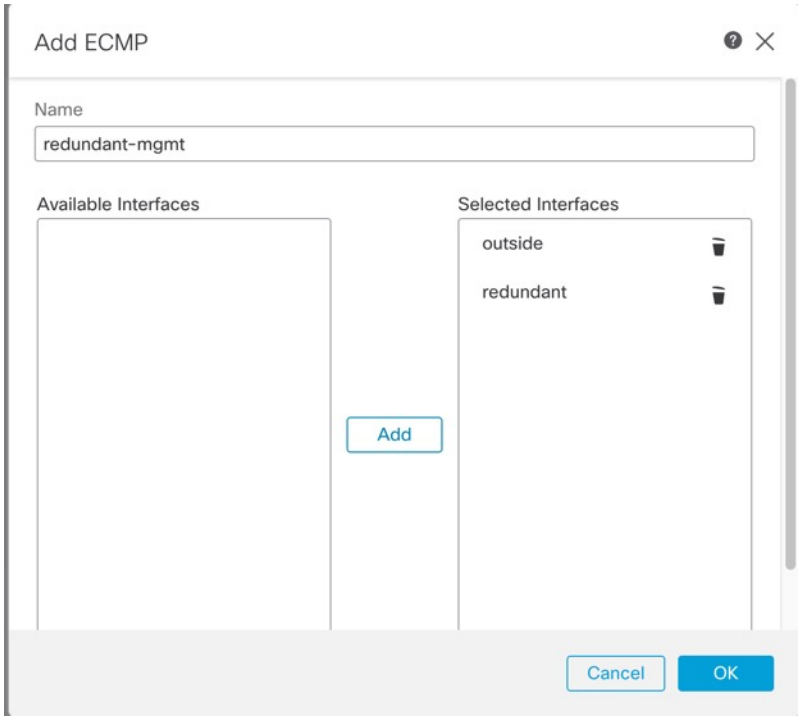


d) Click **Save**.

Step 4 Create an ECMP zone with both interfaces.

a) Click **Routing**.
b) From the virtual router drop-down, choose the virtual router in which the primary and secondary interfaces reside.
c) Click **ECMP**, and then click **Add**.
d) Enter a **Name** for the ECMP zone.
e) Select the primary and secondary interfaces under the **Available Interfaces** box, and then click **Add**.

Figure 291: Add an ECMP zone

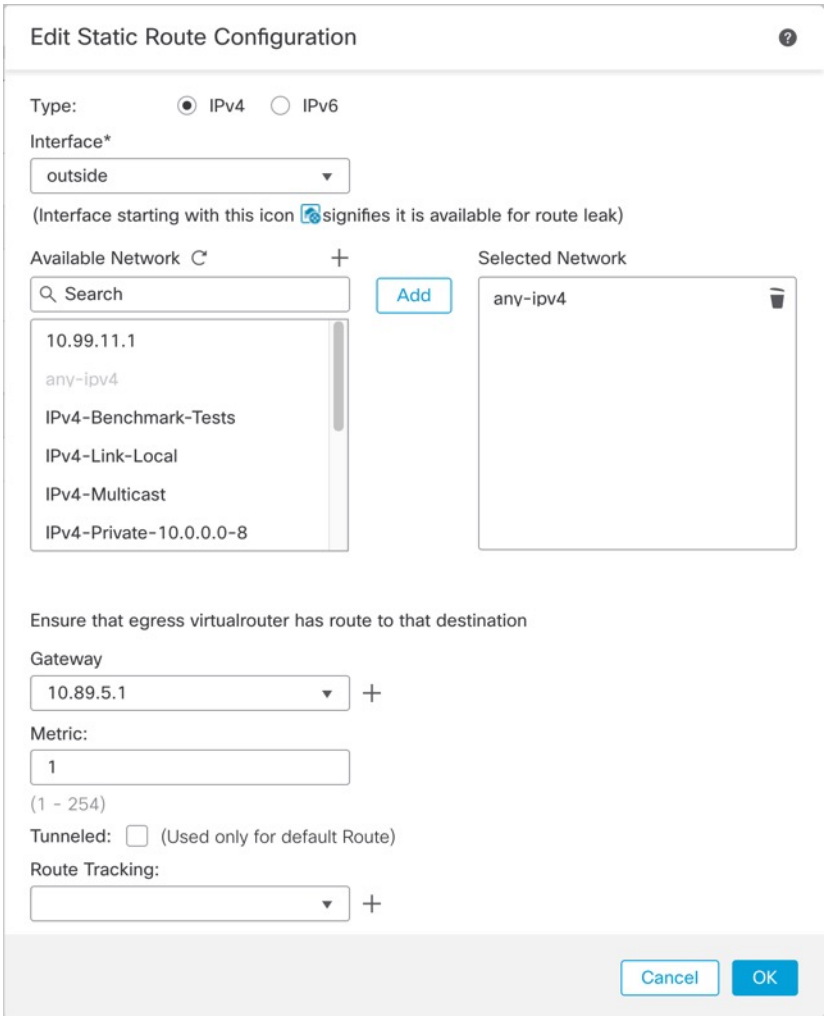


f) Click **OK**, and then **Save**.

Step 5 Add equal-cost default static routes for both interfaces and enable SLA tracking on both.

The routes should be identical except for the gateway and should both have metric 1. The primary interface should already have a default route that you can edit.

Figure 292: Add/Edit static route



a) Click **Static Route**.

b) Either click **Add Route** to add a new route, or click **Edit** (✎) for an existing route.

c) From the **Interface** drop-down, choose the interface.

d) For the destination network, select **any-ipv4** from the **Available Networks** box and click **Add**.

e) Enter the default **Gateway**.

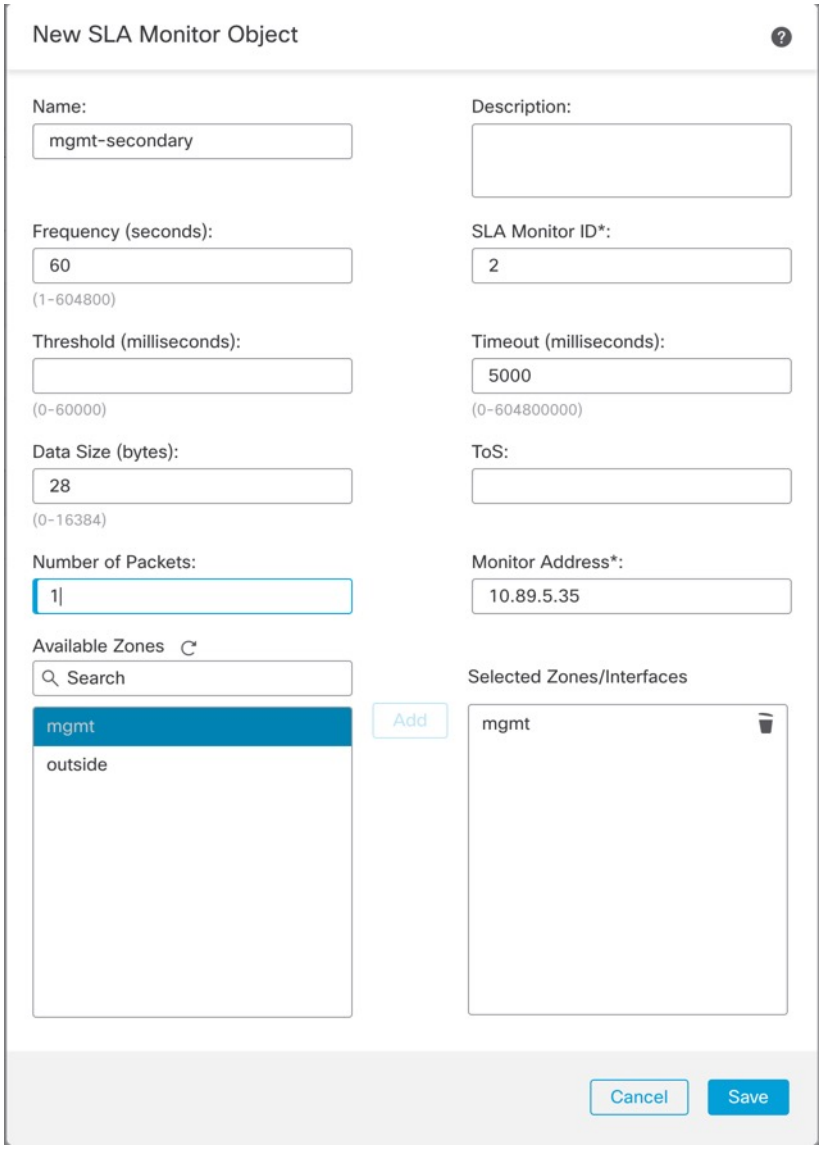
f) For **Route Tracking**, click **Add** (+) to add a new SLA monitor object.

g) Enter the required parameters including the following:

- The **Monitor Address** as the Cloud-Delivered Firewall Management Center IP address.
- The zone for the primary or secondary management interface in **Available Zones**; for example, choose the outside zone for the primary interface object, and the mgmt zone for the secondary interface object.

See Configure an SLA monitor, on page 3034 for more information.

Figure 293: Add SLA monitor



h) Click **Save**, then choose the SLA object you just created in the **Route Tracking** drop-down list.

i) Click **OK**, and then **Save**.

Repeat for the default route for the other management interface.

Step 6 Deploy configuration changes.

To deploy this feature, the Cloud-Delivered Firewall Management Center enables the secondary interface for management traffic, including auto-generated policy-based routing configuration for management traffic to get to the right data interface. The Cloud-Delivered Firewall Management Center also deploys a second instance of the **configure network management-data-interface** command. Note that if you edit the secondary

interface at the CLI, you cannot configure the gateway or otherwise alter the default route, because the static route for this interface can only be edited in the Cloud-Delivered Firewall Management Center.

Change manager access interface settings

Changing manager interface settings on the device or on the Cloud-Delivered Firewall Management Center can disrupt the management connection. Review these scenarios that change interface settings and reestablish the management connection.

Change the device IP address

Change the device IP address, and then update the address in the Cloud-Delivered Firewall Management Center.

Set the device IP address

Use one of these methods to set the manager access interface IP address.

Modify Firewall Threat Defense management interface settings at the CLI

Many management interface settings are set during device initial setup; this task allows you to update them, add event-only interfaces (if supported), or configure static routes, DNS domains, and other advanced settings. Use the CLI when you need direct device access or troubleshooting. If you wish to configure network settings for a data interface used for management, use the device management center interface instead.

Note This topic applies to the dedicated management interface. You can alternatively configure a data interface for management. If you want to change network settings for that interface, you should do so within Cloud-Delivered Firewall Management Center and not at the CLI. If you need to troubleshoot a disrupted management connection, and need to make changes directly on the Firewall Threat Defense , see Modify the Firewall Threat Defense data interface used for management at the CLI, on page 1163 .

For information about the Firewall Threat Defense CLI, see the *Cisco Secure Firewall Threat Defense Command Reference* .

Before you begin

- You can create user accounts that can log into the CLI using the **configure user add** command ; see Add an internal user at the CLI, on page 183 . You can also configure AAA users according to Configure external authentication, on page 1249 .
- Be cautious when making changes to the management interface using SSH. If you cannot re-connect because of a configuration error, access the device console port.
- In high-availability configurations, when you modify the management IP address of a registered device from the device CLI or from the Cloud-Delivered Firewall Management Center , the standby Cloud-Delivered Firewall Management Center does not reflect the changes even after a high-availability synchronization. To ensure that the standby Cloud-Delivered Firewall Management Center is also updated, switch roles between the two Cloud-Delivered Firewall Management Center s, making the secondary Cloud-Delivered Firewall Management Center as the active unit. Modify the management IP address of

the registered device on the **Device Management** page of the standby Cloud-Delivered Firewall Management Center .

Procedure

Step 1 Connect to the device CLI, either from the console port or using SSH.

Step 2 Log in with the Admin username and password.

Step 3 (Firepower 4100/9300 /Secure Firewall 4200 only) Enable the second management interface as an event-only interface.

configure network management-interface enable management1

configure network management-interface disable-management-channel management1

You always need a management interface for management traffic. If your device has a second management interface, you can enable it for event-only traffic.

(Optional) Disable events for the main management interface using the **configure network management-interface disable-events-channel** command. In either case, the device will try to send events on the event-only interface, and if that interface is down, it will send events on the management interface even if you disable the event channel.

You cannot disable both event and management channels on an interface.

To use a separate event interface, you must enable an event-only interface on the Cloud-Delivered Firewall Management Center that is dedicated exclusively to event traffic. See the *Cisco Secure Firewall Management Center Administration Guide* .

Example:

```
> configure network management-interface enable management1
Configuration updated successfully
> configure network management-interface disable-management-channel management1
Configuration updated successfully
>
```

Step 4 Configure the IP address of the management interface and/or event interface:

If you do not specify the *management_interface* argument, then you change the network settings for the default management interface. When configuring an event interface, be sure to specify the *management_interface* argument. The event interface can be on a separate network from the management interface, or on the same network. If you are connected to the interface you are configuring, you will be disconnected. You can re-connect to the new IP address.

a) Configure the IPv4 address:

- Manual configuration:

  configure network ipv4 manual *ip_address netmask gateway_ip* [*management_interface*]

  Note that the *gateway_ip* in this command is used to create the default route for the device. If you configure an event-only interface, then you must enter the *gateway_ip* as part of the command; however, this entry just configures the default route to the value you specify and does not create a separate static route for the eventing interface. If you are using an event-only interface on a different network from the management interface, we recommend that you set the *gateway_ip* for use with

the management interface, and then create a static route separately for the event-only interface using the **configure network static-routes** command.

Example:

```
> configure network ipv4 manual 10.10.10.45 255.255.255.0 10.10.10.1 management1
Setting IPv4 network configuration.
Network settings changed.
>
```

- DHCP (supported on the default management interface only):

  configure network ipv4 dhcp

b) Configure the IPv6 address:

- Stateless autoconfiguration:

  configure network ipv6 router [*management_interface*]

  Example:

  ```
  > configure network ipv6 router management0
  Setting IPv6 network configuration.
  Network settings changed.
  >
  ```

- Manual configuration:

  configure network ipv6 manual *ip6_address ip6_prefix_length* [*ip6_gateway_ip*] [*management_interface*]

  Note that the *ipv6_gateway_ip* in this command is used to create the default route for the device. If you configure an event-only interface, then you must enter the *ipv6_gateway_ip* as part of the command; however, this entry just configures the default route to the value you specify and does not create a separate static route for the eventing interface. If you are using an event-only interface on a different network from the management interface, we recommend that you set the *ipv6_gateway_ip* for use with the management interface, and then create a static route separately for the event-only interface using the **configure network static-routes** command.

  Example:

  ```
  > configure network ipv6 manual 2001:0DB8:BA98::3210 64 management1
  Setting IPv6 network configuration.
  Network settings changed.
  >
  ```

- DHCPv6 (supported on the default management interface only):

  configure network ipv6 dhcp

c) For IPv6, enable or disable ICMPv6 Echo Replies and Destination Unreachable messages. These messages are enabled by default.

configure network ipv6 destination-unreachable { **enable** | **disable** }

configure network ipv6 echo-reply { **enable** | **disable** }

You might want to disable these packets to guard against potential denial of service attacks. Disabling Echo Reply packets means you cannot use IPv6 ping to the device management interfaces for testing purposes.

Example:

```
> configure network ipv6 destination-unreachable disable
> configure network ipv6 echo-reply disable
```

Step 5 Enable a DHCP server on the default management interface to provide IP addresses to connected hosts:

configure network ipv4 dhcp-server-enable *start_ip_address end_ip_address*

Example:

```
> configure network ipv4 dhcp-server-enable 10.10.10.200 10.10.10.254
DHCP Server Enabled
>
```

You can only configure a DHCP server when you set the management interface IP address manually. This command is not supported on the Firewall Management Center Virtual . To display the status of the DHCP server, enter **show network-dhcp-server** :

```
> show network-dhcp-server
DHCP Server Enabled
10.10.10.200-10.10.10.254
```

Step 6 Add a static route for the event-only interface if the Cloud-Delivered Firewall Management Center is on a remote network. Otherwise, all traffic will match the default route through the management interface.

configure network static-routes { **ipv4** | **ipv6** } **add** *management_interface destination_ip netmask_or_prefix gateway_ip*

For the default route, do not use this command; you can only change the default route gateway IP address when you use the **configure network ipv4** or **ipv6** commands (see Step 4, on page 1159).

Example:

```
> configure network static-routes ipv4 add management1 192.168.6.0 255.255.255.0 10.10.10.1
Configuration updated successfully
> configure network static-routes ipv6 add management1 2001:0DB8:AA89::5110 64
2001:0DB8:BA98::3211
Configuration updated successfully
>
```

To display static routes, enter **show network-static-routes** (the default route is not shown):

```
> show network-static-routes
---------------[ IPv4 Static Routes ]---------------
Interface                : management1
Destination              : 192.168.6.0
Gateway                  : 10.10.10.1
Netmask                  : 255.255.255.0
[…]
```

Step 7 Do these settings:

a) Set the hostname:

configure network hostname *name*

Example:

```
> configure network hostname farscape1.cisco.com
```

Syslog messages do not reflect a new hostname until after a reboot.

b) Set the search domains:

configure network dns searchdomains *domain_list*

Example:

```
> configure network dns searchdomains example.com,cisco.com
```

Set the search domain(s) for the device, separated by commas. These domains are added to hostnames when you do not specify a fully-qualified domain name in a command, for example, **ping system** . The domains are used only on the management interface, or for commands that go through the management interface.

c) Set up to 3 DNS servers, separated by commas:

configure network dns servers *dns_ip_list*

Example:

```
> configure network dns servers 10.10.6.5,10.20.89.2,10.80.54.3
```

d) Set the remote management port for communication with the Cloud-Delivered Firewall Management Center :

configure network management-interface tcpport *number*

Example:

```
> configure network management-interface tcpport 8555
```

The Cloud-Delivered Firewall Management Center and managed devices communicate using a two-way, TLS-1.3-encrypted communication channel, which by default is on port 8305. Do not change the management port when using multi-instance mode; only port 8305 is supported.

Note
We recommend that you keep the default settings for the remote management port, but if the management port conflicts with other communications on your network, you can choose a different port. If you change the management port, you must change it for *all* devices in your deployment that need to communicate with each other.

e) (Firewall Threat Defense only) Set the management or eventing interface MTU. The MTU is 1500 bytes by default.

configure network mtu [*bytes*] [*interface_id*]

- *bytes* —Sets the MTU in bytes. For the management interface, the value can be between 64 and 1500 if you enable IPv4, and 1280 to 1500 if you enable IPv6. For the eventing interface, the value can be between 64 and 9000 if you enable IPv4, and 1280 to 9000 if you enable IPv6. If you enable both IPv4 and IPv6, then the minimum is 1280. If you do not enter the *bytes* , you are prompted for a value.
- *interface_id* —Specifies the interface ID on which to set the MTU. Use the **show network** command to see available interface IDs, for example management0, management1, br1, and eth0, depending on the platform. If you do not specify an interface, then the management interface is used.

Example:

```
> configure network mtu 8192 management1
MTU set successfully to 1500 from 8192 for management1
Refreshing Network Config...
NetworkSettings::refreshNetworkConfig MTU value at start 8192
Interface management1 speed is set to '10000baseT/Full'
NetworkSettings::refreshNetworkConfig MTU value at end 8192
>
```

f) Configure an HTTP proxy. The device is configured to directly connect to the internet on ports TCP/443 (HTTPS) and TCP/80 (HTTP). You can use a proxy server, to which you can authenticate via HTTP

Digest. After issuing the command, you are prompted for the HTTP proxy address and port, whether proxy authentication is required, and if it is required, the proxy username, proxy password, and confirmation of the proxy password.

Note
For proxy password on Firewall Threat Defense , you can use A-Z, a-z, and 0-9 characters only.

configure network http-proxy

Example:

```
> configure network http-proxy
Manual proxy configuration
Enter HTTP Proxy address: 10.100.10.10
Enter HTTP Proxy Port: 80
Use Proxy Authentication? (y/n) [n]: Y
Enter Proxy Username: proxyuser
Enter Proxy Password: proxypassword
Confirm Proxy Password: proxypassword
```

Step 8 If you change the device management IP address, then see the following tasks for Cloud-Delivered Firewall Management Center connectivity. Perform these tasks depending on how you identified the Cloud-Delivered Firewall Management Center during initial device setup using the **configure manager add** command :

- **IP address — No action.** If you identified the Cloud-Delivered Firewall Management Center using a reachable IP address, then the management connection will be reestablished automatically after several minutes. We recommend that you also change the device IP address shown in Cloud-Delivered Firewall Management Center to keep the information in sync; see Update the Hostname or IP Address in the Cloud-Delivered Firewall Management Center, on page 1167. This action can help reestablish the connection faster. If you specified an unreachable Cloud-Delivered Firewall Management Center IP address, then you must manually reestablish the connection using Update the Hostname or IP Address in the Cloud-Delivered Firewall Management Center, on page 1167
- **NAT ID only — Manually reestablish the connection.** If you identified the Cloud-Delivered Firewall Management Center using only the NAT ID, then the connection does not automatically reestablish. In this case, change the device management IP address in Cloud-Delivered Firewall Management Center according to Update the Hostname or IP Address in the Cloud-Delivered Firewall Management Center, on page 1167 .

Modify the Firewall Threat Defense data interface used for management at the CLI

This task allows you to specify a new data interface to replace the old interface when the management connection has been disrupted and you need to restore connectivity between the Firewall Threat Defense and the Cloud-Delivered Firewall Management Center.

If the management connection is active, make changes to an existing data interface using the Cloud-Delivered Firewall Management Center (see Modify the Firewall Threat Defense data interface used for management in the GUI, on page 1166). For initial setup of the data management interface, see the **configure network management-data-interface** command.

For high-availability pairs, perform all CLI steps on both units. Within the Cloud-Delivered Firewall Management Center, perform steps only on the active unit. Once the configuration changes are deployed, the standby unit synchronizes configuration and other state information from the active unit.

Note This topic applies to the data interface that you configured for Management, not the dedicated Management interface. If you want to change network settings for the Management interface, see Modify Firewall Threat Defense management interface settings at the CLI , on page 1158.

For information about the Firewall Threat Defense CLI, see the *Cisco Secure Firewall Threat Defense Command Reference*.

Before you begin

If you are changing the data management interface to a new interface, move the current interface cable to the new interface.

Procedure

Step 1 Connect to the device CLI.

Use the console port when executing these commands. If you are performing initial setup, then you may be disconnected from the Management interface. If you are editing the configuration due to a disrupted management connection and have SSH access to the dedicated management interface, you can use the SSH connection.

a) Log in with the **admin** username and password.

Step 2 Disable the interface so you can reconfigure its settings.

configure network management-data-interface disable

Note
If you only want to set a new IPv4 address on the same interface and not make any other changes, you can skip this step. Other changes require you to disable the interface first.

Example:

```
> configure network management-data-interface disable

 Configuration updated successfully..!!


Configuration disable was successful, please update the default route to point to a gateway
 on management interface using the command 'configure network'
```

Step 3 Configure the new data interface for manager access.

configure network management-data-interface

You are then prompted to configure basic network settings for the data interface.

If you change the data management interface to a new interface on the same network, use the same settings as for the previous interface except the interface ID. In addition, for the **Do you wish to clear all the device configuration before applying ? (y/n) [n]:** option, choose **y**. This choice will clear the old data management interface configuration, so that you can successfully reuse the IP address and interface name on the new interface.

```
> configure network management-data-interface
```

```
Data interface to use for management: ethernet1/4
Specify a name for the interface [outside]: internet
IP address (manual / dhcp) [dhcp]: manual
IPv4/IPv6 address: 10.10.6.7
Netmask/IPv6 Prefix: 255.255.255.0
Default Gateway: 10.10.6.1
Comma-separated list of DNS servers [none]: 208.67.222.222,208.67.220.220
DDNS server update URL [none]:
Do you wish to clear all the device configuration before applying ? (y/n) [n]: y

Configuration done with option to allow manager access from any network, if you wish to
change the manager access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.

>
```

a) Limit data interface access to the Cloud-Delivered Firewall Management Center on a specific network.

configure network management-data-interface client *ip_address netmask*

By default, all networks are allowed.

Step 4 Update the Hostname or IP Address in the Cloud-Delivered Firewall Management Center, on page 1167.

The connection may be reestablished automatically, but disabling and reenabling the connection in the Cloud-Delivered Firewall Management Center will help the connection reestablish faster. Or you may need to update the device IP address in the Cloud-Delivered Firewall Management Center according to the linked procedure.

Step 5 Check that the management connection was reestablished.

sftunnel-status-brief

See the sample output for a connection that is up, with peer channel and heartbeat information:

```
> sftunnel-status-brief
PEER:10.10.17.202
Peer channel Channel-A is valid type (CONTROL), using 'eth0', connected to '10.10.17.202'
via '10.10.17.222'
Peer channel Channel-B is valid type (EVENT), using 'eth0', connected to '10.10.17.202' via
 '10.10.17.222'
Registration: Completed.
IPv4 Connection to peer '10.10.17.202' Start Time: Wed Jun 10 14:27:12 2020 UTC
Heartbeat Send Time: Mon Jun 15 09:02:08 2020 UTC
Heartbeat Received Time: Mon Jun 15 09:02:16 2020 UTC
```

Step 6 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and click **Edit** (). In the **Device** area, against the **Management** field, click **Refresh** next to **Manager Access - Configuration Details**.

The Cloud-Delivered Firewall Management Center detects the interface and default route configuration changes and blocks deployment to the device. When you change the data interface settings locally on the device, you must reconcile those changes in the Cloud-Delivered Firewall Management Center manually. You can view the discrepancies between the Cloud-Delivered Firewall Management Center and the device on the **Configuration** tab.

Step 7 Choose **Interfaces**, and make the changes.

a) Remove the IP address and name from the old data management interface and disable manager access for this interface.
b) Configure the new data management interface with the new settings (the ones you used at the CLI) and enable manager access for it.

Step 8 Under the **Routing** tab, click **Static Route**, and change the default route from the old data management interface to the new one.

Step 9 Return to the **Manager Access - Configuration Details** dialog box, and click **Acknowledge** to remove the deployment block.

The next time you deploy, the Cloud-Delivered Firewall Management Center configuration will overwrite any remaining conflicting settings on the Firewall Threat Defense. It is your responsibility to manually fix the configuration in the Cloud-Delivered Firewall Management Center before you re-deploy.

You will see expected messages of "Config was cleared" and "Manager access changed and acknowledged."

Modify the Firewall Threat Defense data interface used for management in the GUI

Configure the data interface IP address settings to change the management connection or update the interface used for manager access.

If the management connection is up but you want to change the IP address of the data interface used for manager access, follow these steps. For example, if you register a device using zero-touch provisioning, then you need to change the IP address to a static address before you can enable high availability.

You can alternatively change interface settings at the CLI, but we recommend using that method only if the management connection is down. Any changes you make at the CLI should be replicated in the GUI.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) next to the device.

Step 2 Choose **Interfaces**.

Step 3 If you want to change the interface used for manager access:

a) Remove the IP address and name from the old data management interface and disable manager access for this interface.
b) Configure the new data management interface with the new settings and enable manager access for it.
c) If you use a static IP address, you are reminded to make sure you have a default route. Click **Yes**.
d) Click **OK** to exit the interface.
e) Click **Save** on the **Interfaces** page.

Step 4 If you only want to change the IP address:

a) Change the IP address.
b) For a static IP address, you are reminded to make sure you have a default route. Click **Yes**.
c) Click **OK** to exit the interface.
d) Click **Save** on the **Interfaces** page.

Step 5 Click the **Routing** tab, click **Static Route**, and then add or change the default or static route for the manager access interface.

Step 6 Deploy configuration changes.

The Cloud-Delivered Firewall Management Center will deploy the configuration changes over the current connection. After the deployment, the data interface will have a new IP address, so the management connection will need to be reestablished.

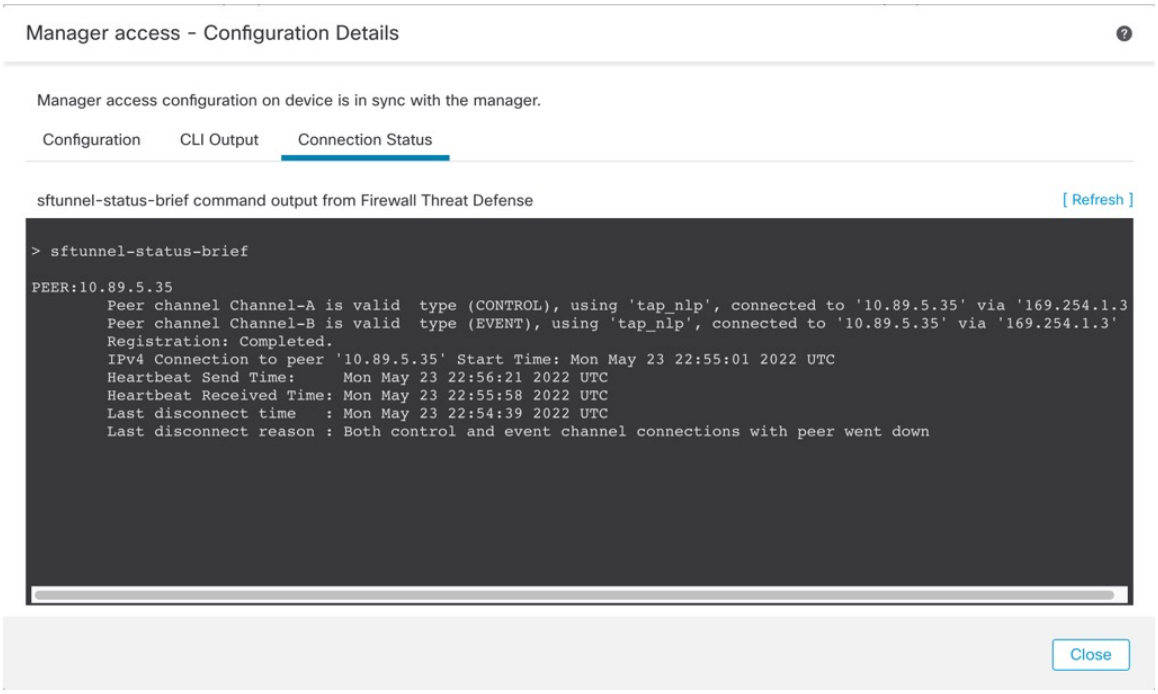
Step 7 Update the Hostname or IP Address in the Cloud-Delivered Firewall Management Center, on page 1167.

Step 8 Ensure the management connection is reestablished.

In the **Device** area, against the **Management** field, click **Manager Access Details: Configuration** and then click **Connection Status**.

The following status shows a successful connection for a data interface, showing the internal "tap_nlp" interface.

Figure 294: Connection status



If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 1183.

Update the Hostname or IP Address in the Cloud-Delivered Firewall Management Center

This task allows you to manually update the hostname or IP address on the managing Cloud-Delivered Firewall Management Center when device network settings have been modified outside of the management interface.

If you edit the hostname or IP address of a device after you added it to the Cloud-Delivered Firewall Management Center (using the device's CLI, for example), you need to use this procedure to manually update the hostname or IP address on the managing Cloud-Delivered Firewall Management Center.

To change the device management IP address on the device, see Modify Firewall Threat Defense management interface settings at the CLI , on page 1158.

If you used only the NAT ID when registering the device, then the IP shows as **NO-IP** on this page, and you do not need to update the IP address/hostname.

If you used zero-touch provisioning to register the device on the outside interface, the hostname is automatically generated along with a matching DDNS configuration. You cannot edit the hostname in this case.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to modify management options, click **Edit** ( ).

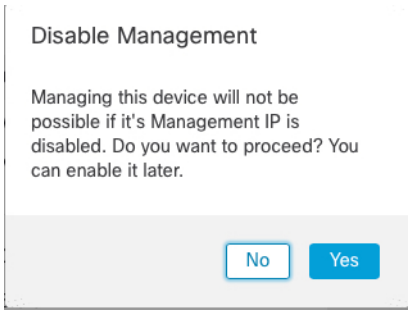
Step 3 Click **Device** and view the **Management** area.

Step 4 Disable management temporarily by clicking the slider so it is disabled **Slider disabled** ( ).

Figure 295: Disable management



You are prompted to proceed with disabling management; click **Yes**.



Disabling management blocks the connection between the Cloud-Delivered Firewall Management Center and the device, but does **not** unregister the device from the Cloud-Delivered Firewall Management Center.

Step 5 Edit the **Remote Host Address** IP address and optional **Secondary Address** (when using a redundant data interface) or hostname by clicking **Edit** ( ).

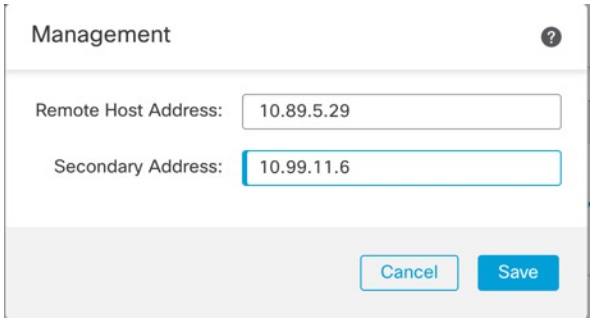
Figure 296: Edit management address



Step 6 In the **Management** dialog box, modify the name or IP address in the **Remote Host Address** field and the optional **Secondary Address** field. Click **Save**.

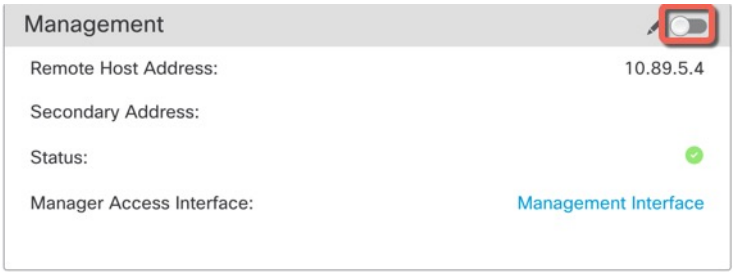
For information about using a secondary manager access data interface, see Configure a redundant manager access data interface, on page 1153.

Figure 297: Management IP address



Step 7 Reenable management by clicking the slider so it is enabled **Slider enabled** ( ).

Figure 298: Enable management connection



Change the Cloud-Delivered Firewall Management Center IP Address

This task ensures the Cloud-Delivered Firewall Management Center IP address or hostname matches between the management center and device CLI, maintaining consistent configurations and network resiliency.

If you change the Cloud-Delivered Firewall Management Center IP address or hostname, you should also change the value at the device CLI so the configurations match. Although in most cases, the management connection will be reestablished without changing the Cloud-Delivered Firewall Management Center IP

address or hostname on the device, in at least one case, you must perform this task for the connection to be reestablished: when you added the device to the Cloud-Delivered Firewall Management Center and you specified the NAT ID only. Even in other cases, we recommend keeping the Cloud-Delivered Firewall Management Center IP address or hostname up to date for extra network resiliency.

Procedure

Step 1 Change the Cloud-Delivered Firewall Management Center IP address.

Caution
Be careful when making changes to the Cloud-Delivered Firewall Management Center interface to which you are connected; if you cannot re-connect because of a configuration error, you need to access the Cloud-Delivered Firewall Management Center console port to re-configure the network settings in the Linux shell. You must contact Cisco TAC to guide you in this operation.

a) Choose **System** (⚙) > **Configuration** > **Management Interfaces**.
b) In the **Interfaces** area, click **Edit** next to the interface that you want to configure.
c) Change the IP address, and click **Save**.

Step 2 At the Firewall Threat Defense CLI, view the Cloud-Delivered Firewall Management Center identifier.

show managers

Example:

```
> show managers
Type                      : Manager
Host                      : 10.10.1.4
Display name              : 10.10.1.4
Identifier                : f7ffad78-bf16-11ec-a737-baa2f76ef602
Registration              : Completed
Management type           : Configuration
```

Step 3 At the Firewall Threat Defense CLI, edit the Cloud-Delivered Firewall Management Center IP address or hostname.

configure manager edit *identifier* {**hostname** {*ip_address* | *hostname*} | **displayname** *display_name*}

If the Cloud-Delivered Firewall Management Center was originally identified by **DONTRESOLVE** and a NAT ID, you can change the value to a hostname or IP address using this command. You cannot change an IP address or hostname to **DONTRESOLVE**.

The management connection will go down, and then reestablish. You can monitor the state of the connection using the **sftunnel-status** command.

Example:

```
> configure manager edit f7ffad78-bf16-11ec-a737-baa2f76ef602 hostname 10.10.5.1
```

Change both Cloud-Delivered Firewall Management Center and Threat Defense IP addresses

You might want to change both Cloud-Delivered Firewall Management Center and Firewall Threat Defense IP addresses if you need to move them to a new network.

Procedure

Step 1 Disable the management connection.

For a high-availability pair or cluster, perform these steps on all units.

a) Choose **Devices** > **Device Management**.

b) Next to the device, click **Edit** ( ).

c) Click **Device**, and view the **Management** area.

d) Disable management temporarily by clicking the slider so it is disabled ( ).

Figure 299: Disable Management



You are prompted to proceed with disabling management; click **Yes**.

Step 2 Change the device IP address in the Cloud-Delivered Firewall Management Center to the new device IP address.

You will change the IP address on the device later.

For a high-availability pair or cluster, perform these steps on all units.

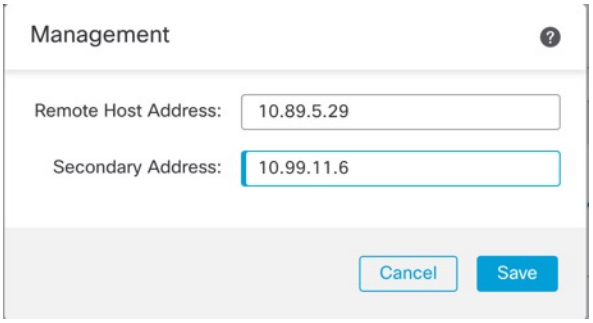
a) Edit the **Remote Host Address** IP address and optional **Secondary Address** (when using a redundant data interface) or hostname by clicking **Edit** ( ).

Figure 300: Edit Management Address



b) In the **Management** dialog box, modify the name or IP address in the **Remote Host Address** field and the optional **Secondary Address** field, and click **Save**.

Figure 301: Management IP Address



Step 3 Change the Cloud-Delivered Firewall Management Center IP address.

Caution
Be careful when making changes to the Cloud-Delivered Firewall Management Center interface to which you are connected. If you cannot re-connect because of a configuration error, you need to access the Cloud-Delivered Firewall Management Center console port to re-configure the network settings in the Linux shell. You must contact Cisco TAC to guide you in this operation.

a) Choose **System** (⚙) > **Configuration** > **Management Interfaces**.
b) In the **Interfaces** area, click **Edit** next to the interface that you want to configure.
c) Change the IP address, and click **Save**.

Step 4 Change the manager IP address on the device.

For a high-availability pair or cluster, perform these steps on all units.

a) At the Firewall Threat Defense CLI, view the Cloud-Delivered Firewall Management Center identifier.

show managers

Example:

```
> show managers
Type                      : Manager
Host                      : 10.10.1.4
Display name              : 10.10.1.4
Identifier                : f7ffad78-bf16-11ec-a737-baa2f76ef602
Registration              : Completed
Management type           : Configuration
```

b) Edit the Cloud-Delivered Firewall Management Center IP address or hostname.

configure manager edit *identifier* { **hostname** { *ip_address* | *hostname* } | **displayname** *display_name* }

If the Cloud-Delivered Firewall Management Center was originally identified by **DONTRESOLVE** and a NAT ID, you can change the value to a hostname or IP address using this command. You cannot change an IP address or hostname to **DONTRESOLVE**.

Example:

```
                                > configure manager edit
f7ffad78-bf16-11ec-a737-baa2f76ef602
```

```
hostname
10.10.5.1
```

Step 5 Change the IP address of the manager access interface at the console port.

For a high-availability pair or cluster, perform these steps on all units.

If you use the dedicated Management interface:

configure network ipv4

configure network ipv6

If you use the dedicated Management interface:

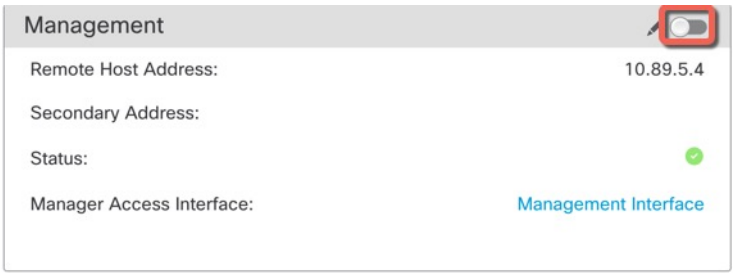
configure network management-data-interface disable

configure network management-data-interface

Step 6 Re-enable management by clicking the slider ( ).

For a high-availability pair or cluster, perform these steps on all units.

Figure 302: Enable Management Connection



Step 7 (If you use a data interface for manager access) Refresh the data interface settings in the Cloud-Delivered Firewall Management Center.

For a high-availability pair, perform this step on both units.

a) Choose **Devices** > **Device Management**, and click **Manager Access - Configuration Details** , and then click **Refresh**.
b) Choose **Devices** > **Device Management**, and click the **Interfaces** tab and set the IP address to match the new address.
c) Return to the **Manager Access - Configuration Details** dialog box, and click **Acknowledge** to remove the deployment block.

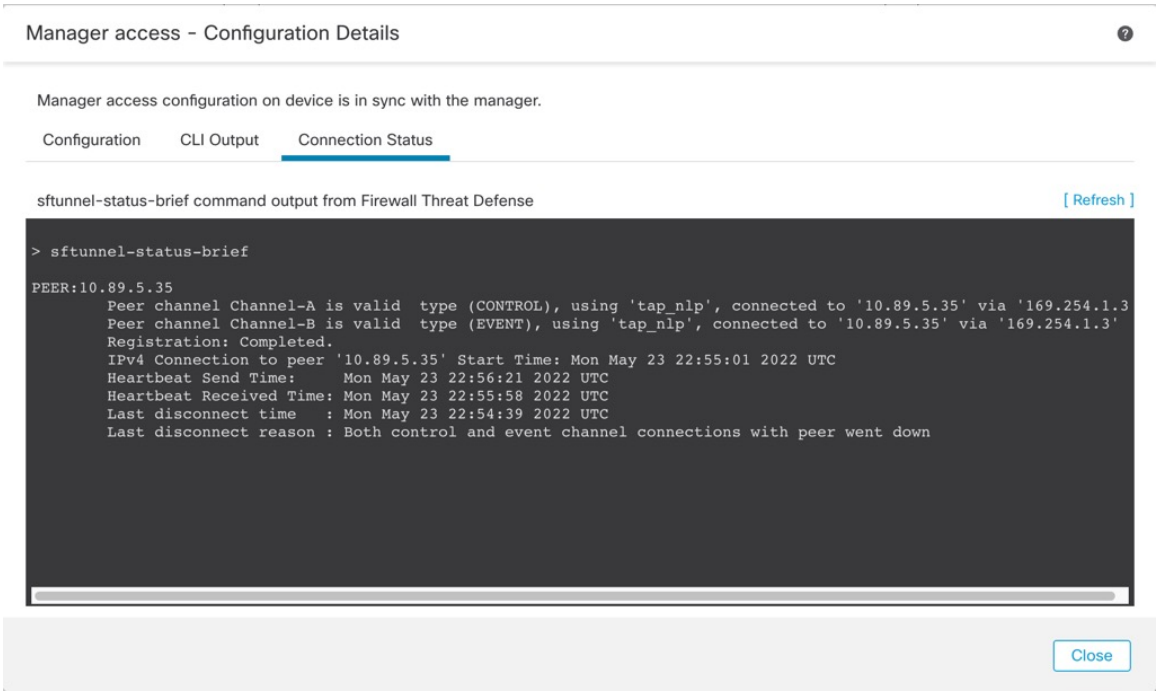
Step 8 Ensure the management connection is reestablished.

In the Cloud-Delivered Firewall Management Center, check the management connection status. Navigate to the **Devices** > **Device Management**, and click **Management** section under the **Device** tab. Then, click **Manager Access - Configuration Details** to view the **Connection Status** page.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

This status indicates a successful connection for a data interface and displays the internal "tap_nlp" interface.

Figure 303: Connection Status



Step 9 (For a high-availability Cloud-Delivered Firewall Management Center pair) Repeat configuration changes on the secondary Cloud-Delivered Firewall Management Center.

a) Change the secondary Cloud-Delivered Firewall Management Center IP address.
b) Specify the new peer addresses on both units.
c) Make the secondary unit the active unit.
d) Disable the device management connection.
e) Change the device IP address in the Cloud-Delivered Firewall Management Center.
f) Re-enable the management connection.

Change the manager access interface

After you register the device, you can change the manager access interface between the Management interface and a data interface.

Change the manager access interface from management to data

You can manage the Firewall Threat Defense from either the dedicated management interface or from a data interface. If you want to change the manager access interface after you added the device to the Cloud-Delivered Firewall Management Center , follow these steps to migrate from the Management interface to a data interface. To migrate the other direction, refer to Change the manager access interface from data to management, on page 1178 .

When you initiate the manager access migration from management to data, the Cloud-Delivered Firewall Management Center applies a block on deployment to the Firewall Threat Defense. To remove the block, enable manager access on the data interface.

Before you begin

For high-availability pairs:

- Ensure that you configure a standby IP address for data-managed interfaces before changing the manager access interface from management to data. For more information about configuring standby IP address, refer to Configure standby IP addresses and interface monitoring, on page 708.
- For high-availability pairs, unless stated otherwise, perform all steps only on the active unit. After deploying the configuration changes, the standby unit synchronizes configuration and other state information from the active unit.
-

For high-availability pairs, unless stated otherwise, perform all steps only on the active unit. After deploying the configuration changes, the standby unit synchronizes configuration and other state information from the active unit.

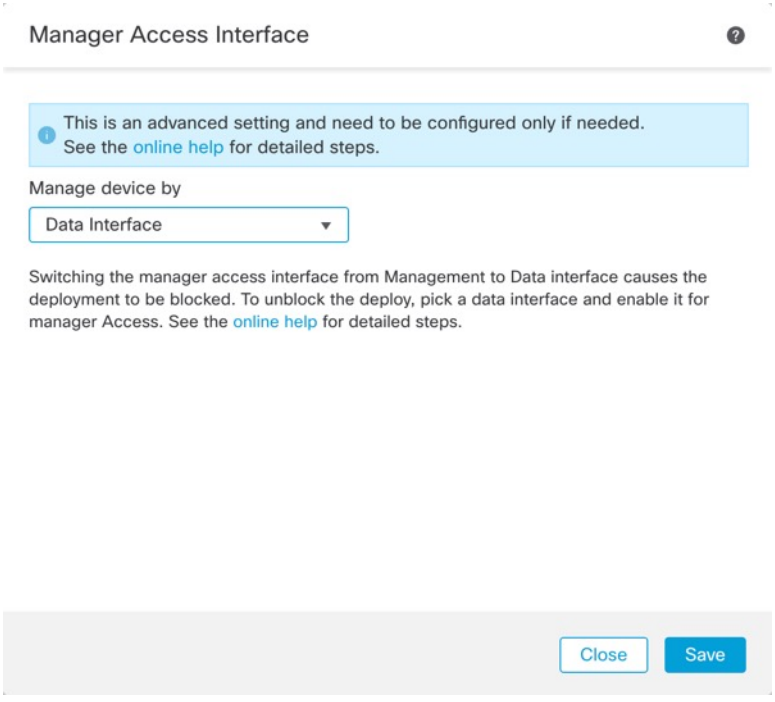
Procedure

Step 1 Initiate the interface migration.

a) On the **Devices** > **Device Management** page, click **Edit** (✎) for the device. Click **Device** , and in the **Management** area, click the **Manager Access Interface** link.

The **Manager Access Interface** field shows the current Management interface. When you click the link, choose the new interface type, **Data Interface** , in the **Manage device by** drop-down list.

Figure 304: Manager Access Interface



b) Click **OK** and then **Close** .

You must now complete the remaining steps in this procedure to enable manager access on the data interface. The **Management** area now shows **Manager Access Interface: Data Interface** , and **Manager Access Details: Configuration** .

Figure 305: Manager Access



If you click **Configuration** , the **Manager Access - Configuration Details** dialog box opens. The **Manager Access Mode** shows a Deploy pending state.

Step 2 Enable manager access on the data interface (s) . Click **Interfaces** , click **Edit** ( ) for the interface, and then click **Manager Access** .

Check **Enable management access** and click **OK** . By default, all networks are allowed, but you can limit access as long as the Cloud-Delivered Firewall Management Center address is allowed.

If the manager access interface uses a static IP address, you are reminded to configure routing for it.

Click **Save** on the **Interfaces** page. See Configure routed mode interfaces, on page 1080 for more information about interface settings. You can enable manager access on one routed data interface , plus an optional secondary interface . Make sure these interfaces are fully configured with a name and IP address and that they are enabled.

If you use a secondary interface for redundancy, see Configure a redundant manager access data interface, on page 1153 for additional required configuration.

Step 3 (Optional) If you use DHCP for the interface, enable the web type DDNS method on the DDNS page. Navigate to **Devices** > **Device Management** , and then click **DDNS** under the **DHCP** tab.

See Configure Dynamic DNS, on page 1234 . DDNS ensures the Cloud-Delivered Firewall Management Center can reach the Firewall Threat Defense at its Fully-Qualified Domain Name (FQDN) if the FTD's IP address changes.

Step 4 Make sure the Firewall Threat Defense can route to the Cloud-Delivered Firewall Management Center through the data interface. Add a static route if necessary on the Static Route page. Navigate to **Devices** > **Device Management** and then click **Static Route** under the **Routing** tab.

See Add a static route, on page 1361 .

Step 5 (Optional) Configure DNS in a Platform Settings policy: choose **Devices** > **Platform Settings** , and click **DNS** . Apply the policy to this device.

See Configure DNS server settings, on page 1246 . DNS is required if you use DDNS. You may also use DNS for FQDNs in your security policies.

Step 6 (Optional) Enable SSH for the data interface in a Platform Settings policy, and apply it to this device at **Devices** > **Device Management** page. Click **Edit** ( ) for the device and then click **SSH Access** .

See Configure Secure Shell (SSH) access, on page 1262 . SSH is not enabled by default on the data interfaces, so if you want to manage the Firewall Threat Defense using SSH, you need to explicitly allow it.

Step 7 Deploy configuration changes.

You will see a validation error to confirm that you are changing the manager access interface. Check **Ignore warnings** and deploy again.

The Cloud-Delivered Firewall Management Center will deploy the configuration changes over the current Management interface. After the deployment, the data interface is now ready for use, but the original management connection to Management is still active.

Step 8 After deployment, perform these actions.

a) At the Firewall Threat Defense CLI (preferably from the console port), set the Management interface to use a static IP address and set the gateway to use the data interfaces. For high availability, perform this step on both units.

```
configure network
{
ipv4
 |
ipv6
}
manual
ip_address netmask
data-interfaces
```

- *ip_address netmask* —Although you do not plan to use the Management interface, you must set a static IP address, for example, a private address so that you can set the gateway to **data-interfaces** (see the next bullet). You cannot use DHCP because the default route, which must be **data-interfaces** , might be overwritten with one received from the DHCP server.
- **data-interfaces** —This setting forwards management traffic over the backplane so it can be routed through the manager access data interface.

We recommend that you use the console port instead of an SSH connection because when you change the Management interface network settings, your SSH session will be disconnected.

b) If necessary, re-cable the Firewall Threat Defense so it can reach the Cloud-Delivered Firewall Management Center on the data interface. For high availability, perform this step on both units.

c) In the Cloud-Delivered Firewall Management Center, disable the management connection for the Firewall Threat Defense in the **Devices** > **Device Management** page in the **Management** area under the **Device** tab, and then re-enable the connection.

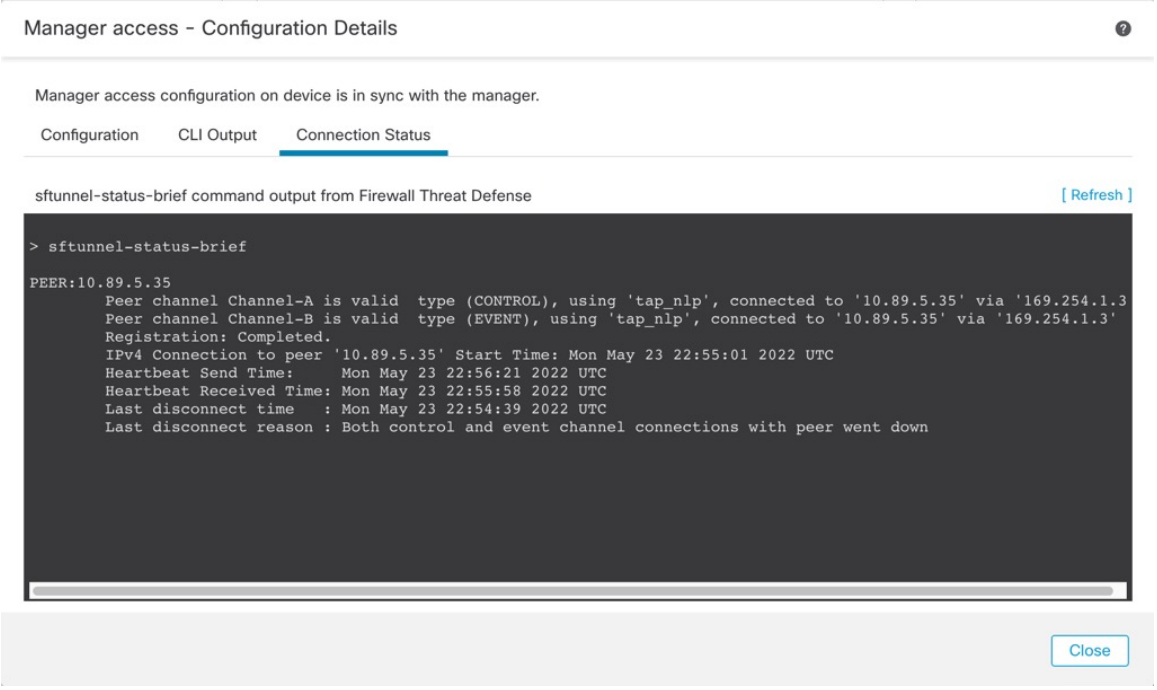
Step 9 Ensure the management connection is reestablished.

In the **Devices** > **Device Management** page, click **Manager Access Details: Configuration** and then click **Connection Status** .

Alternatively, you can check at the Firewall Threat Defense CLI. Enter the **sftunnel-status-brief** command to view the management connection status.

This status indicates a successful connection for a data interface, showing the internal "tap_nlp" interface.

Figure 306: Connection Status



If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 1183 .

Change the manager access interface from data to management

This task enables you to migrate manager access from a data interface to the dedicated management interface, allowing you to manage the Firewall Threat Defense through the Management interface instead of a data interface.

You can manage the Firewall Threat Defense from either the dedicated management interface or from a data interface. If you want to change the manager access interface after you added the device to the Cloud-Delivered Firewall Management Center, follow these steps to migrate from a data interface to the management interface. To migrate the other direction, see Change the manager access interface from management to data, on page 1174.

Initiating the manager access migration from data to management causes the Cloud-Delivered Firewall Management Center to apply a block on deployment to the Firewall Threat Defense. You must disable manager access on the data interface to remove the block.

Before you begin

For high-availability pairs, unless stated otherwise, perform all steps only on the active unit. Once the configuration changes are deployed, the standby unit synchronizes configuration and other state information from the active unit.

Follow these steps to disable manager access on a data interface, and also configure other required settings.

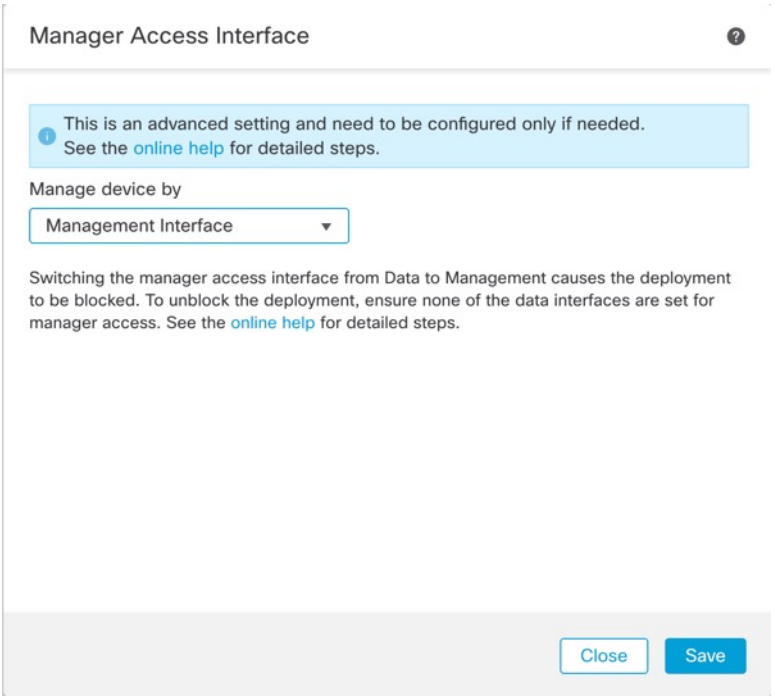
Procedure

Step 1 Initiate the interface migration.

a) On the **Devices** > **Device Management** page, click **Edit** (✎) for the device. Click **Device**, and in the **Management** area, click the link for **Manager Access Interface**.

The **Manager Access Interface** field shows the current management interface as data. When you click the link, choose the new interface type, **Management Interface**, in the **Manage device by** drop-down list.

Figure 307: Manager access interface



b) Click **Save**.

Click **OK** and then **Close**.

You must now complete the remaining steps in this procedure to enable manager access on the Management interface. The **Management** area now shows the **Manager Access Interface: Management Interface**.

Figure 308: Manager access



Step 2 Disable manager access on the data interface(s). Click **Interfaces**, click **Edit** (✎) for the interface, and then click **Manager Access**.

Uncheck **Enable management access** and click **OK**. Click **Save** on the **Interfaces** page. This step removes the block on deployment.

Step 3 If you have not already done so, configure DNS settings for the data interface in a Platform Setting policy, and apply it to this device at **Devices** > **Device Management** page. Click **Edit** (✎) for the device and then click **DNS**.

See Configure DNS server settings, on page 1246. The Cloud-Delivered Firewall Management Center deployment that disables manager access on the data interface will remove any local DNS configuration. If that DNS server is used in any security policy, such as an FQDN in an Access Rule, then you must re-apply the DNS configuration using the Cloud-Delivered Firewall Management Center.

Step 4 Deploy configuration changes.

The Cloud-Delivered Firewall Management Center will deploy the configuration changes over the current data interface.

Step 5 If necessary, re-cable the Firewall Threat Defense so it can reach the Cloud-Delivered Firewall Management Center on the Management interface. For High Availability, perform this step on both units.

Step 6 At the Firewall Threat Defense CLI, configure the Management interface IP address and gateway using a static IP address or DHCP. For high availability, perform this step on both units.

When you originally configured the data interface for manager access, the Management gateway was set to data-interfaces, which forwarded management traffic over the backplane so it could be routed through the manager access data interface. You now need to set an IP address for the gateway on the management network.

Static IP address:

configure network {**ipv4** | **ipv6**} **manual** *ip_address netmask gateway_ip*

DHCP:

configure network {**ipv4** | **ipv6**} **DHCP**

Step 7 In the Cloud-Delivered Firewall Management Center, disable the management connection for the Firewall Threat Defense in the **Devices** > **Device Management**, **Management** section under **Device** tab, and then reenable the connection.

Step 8 Ensure the management connection is reestablished.

In the Cloud-Delivered Firewall Management Center, check the management connection status on the **Devices** > **Device Management**, **Management** section under **Device** tab, **Status** field or view notifications in the Cloud-Delivered Firewall Management Center.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 1183.

Troubleshoot the management connection

This reference provides information to help diagnose and resolve management connection issues.

Manually Roll Back the configuration if the Cloud-Delivered Firewall Management Center loses Connectivity

Use this task to restore management connectivity when a configuration change deployed from the Cloud-Delivered Firewall Management Center affects network connectivity and you need to roll back to the last-deployed configuration.

If you use a data interface on the Firewall Threat Defense for manager access and deploy a configuration change from the Cloud-Delivered Firewall Management Center that affects network connectivity, you can roll back the configuration on the Firewall Threat Defense to the last-deployed configuration to restore management connectivity. Afterward, update the configuration settings in Cloud-Delivered Firewall Management Center to maintain network connectivity, and re-deploy. You can use the rollback feature even if you do not lose connectivity; it is not limited to this troubleshooting situation.

Alternatively, you can enable auto rollback of the configuration if you lose connectivity after a deployment; see Edit deployment settings, on page 1192.

See these guidelines:

- Only the previous deployment is available locally on the Firewall Threat Defense; you cannot roll back to any earlier deployments.
- Rollback is supported for high availability but not supported for clustering deployments.
- Rollback is not supported immediately after high availability creation.
- The rollback only affects configurations that you can set in the Cloud-Delivered Firewall Management Center. For example, the rollback does not affect any local configuration related to the dedicated Management interface, which you can only configure at the Firewall Threat Defense CLI. Note that if you changed data interface settings after the last Cloud-Delivered Firewall Management Center deployment using the **configure network management-data-interface** command, and then you use the rollback command, those settings will not be preserved; they will roll back to the last-deployed Cloud-Delivered Firewall Management Center settings.
- UCAPL/CC mode cannot be rolled back.
- Out-of-band SCEP certificate data that was updated during the previous deployment cannot be rolled back.
- During the rollback, connections will drop because the current configuration will be cleared.

Procedure

Step 1 At the Firewall Threat Defense CLI, roll back to the previous configuration.

configure policy rollback

Note
For a high availability pair, this command is allowed only on the active unit.

After the rollback, the Firewall Threat Defense notifies the Cloud-Delivered Firewall Management Center about the successful completion. The deployment screen In the Cloud-Delivered Firewall Management Center displays a banner indicating that the configuration was rolled back.

Note
If the rollback fails and the Cloud-Delivered Firewall Management Center management is restored, refer to https://www.cisco.com/c/en/us/support/docs/security/firepower-ngfw-virtual/215258-troubleshooting-firepower-threat-defense.html for common deployment problems. In some cases, the rollback may fail after the Cloud-Delivered Firewall Management Center management access is restored. Resolve any Cloud-Delivered Firewall Management Center configuration issues and redeploy from the Cloud-Delivered Firewall Management Center.

Example:

For the Firewall Threat Defense that uses a data interface for manager access:

```
> configure policy rollback

The last deployment to this FTD was on June 1, 2020 and its status was Successful.
Do you want to continue [Y/N]?

Y

Rolling back complete configuration on the FTD. This will take time.
.....................
Policy rollback was successful on the FTD.
Configuration has been reverted back to transaction id:
Following is the rollback summary:
.....................
.....................
>
```

Example:

For Firewall Threat Defenses in a high availability pair that use a data interface for Cloud-Delivered Firewall Management Center access:

```
> configure policy rollback

Checking Eligibility ....
============= DEVICE DETAILS =============
Device Version: 7.2.0
Device Type: FTD
Device Mode: Offbox
Device in HA: true
Is HA disabled: false
HA state: active - standby ready
==========================================
Device is eligible for policy rollback
Do you want to continue [YES/NO]?

YES

Starting rollback...
    Preparing policy configuration on the device.           Status: success
    Applying updated policy configuration on the device.    Status: success
    Applying Lina File Configuration on the device.         Status: success
    Applying Lina Configuration on the device.              Status: success
    Commit Lina Configuration.                              Status: success
    Commit Lina File Configuration.                         Status: success
    Commit Lina File Configuration.                         Status: success
```

```
================================================================
POLICY ROLLBACK STATUS: SUCCESS
================================================================
>
```

Step 2 Check that the management connection was reestablished.

In Cloud-Delivered Firewall Management Center, check the management connection status on the Connection Status page. Navigate to **Devices** > **Device Management** and then navigate to **Management** area under the **Devices** tab. Then in the **Manager Access - Configuration Details** screen, click **Connection Status**.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 1183.

Troubleshoot Management Connectivity on a Data Interface

This topic helps you troubleshoot the loss of management connectivity when using a data interface for manager access instead of the dedicated Management interface.

If you use a data interface for manager access instead of the dedicated Management interface, exercise caution when changing the interface and network settings for the Firewall Threat Defense in the Cloud-Delivered Firewall Management Center to avoid disrupting the connection. If you change the management interface type after adding the Firewall Threat Defense to the Cloud-Delivered Firewall Management Center (switching between data and Management interfaces), misconfiguration of interfaces and network settings can result in loss of management connectivity.

Use the following methods to troubleshoot management connectivity:

- View management connection status: Check the status in the **Devices** > **Device Management** page or use **sftunnel-status-brief** at the CLI.
- View network information: Use **show network** to verify Management and manager access data interface settings.
- Check registration: Use **show managers** to confirm registration completion.
- Ping the manager: Use **ping** or **ping system** to test connectivity.
- Capture packets: Use **capture** on the `nlp_int_tap` interface to trace management traffic.
- Check interface status: Use **show interface detail** for the `nlp_int_tap` interface.
- Check routing and NAT: Use **show route** and **show nat** to verify routing and internal NAT rules.
- Check DDNS: Use **debug ddns** and related commands to verify DDNS updates.

```
> sftunnel-status-brief
PEER:10.10.17.202
Registration: Completed.
Connection to peer '10.10.17.202' Attempted at Mon Jun 15 09:21:57 2020 UTC
Last disconnect time : Mon Jun 15 09:19:09 2020 UTC
Last disconnect reason : Both control and event channel connections with peer went
down
```

```
> show network
===============[ System Information ]===============
Hostname                  : FTD-4
Domains                   : cisco.com
DNS Servers               : 72.163.47.11
DNS from router           : enabled
Management port           : 8305
IPv4 Default route
Gateway                 : data-interfaces
...
```

Note At the Firewall Threat Defense CLI, check that the Cloud-Delivered Firewall Management Center registration was completed. Note that this command will not show the *current* status of the management connection.

Refer to https://cisco.com/go/fmc-reg-error for Cloud-Delivered Firewall Management Center log file information.

Troubleshoot management connectivity on a data interface in a High Availability pair

This topic provides troubleshooting steps for management connectivity issues on data interfaces in High Availability, covering scenarios involving interface connectivity failures and ISP changes.

The management connection between the active peer and the Security Cloud Control can be disrupted due to the following reasons:

- Connectivity issues on the data interface that is used for management on the active unit.
- Change in internet service provider (ISP).

Follow these procedures to restore connectivity based on the identified cause:

- Connectivity issues: Manually fail over to the standby unit and configure a new data interface for Security Cloud Control access.
- ISP change: Manually update the new network details on the active unit using CLI commands to restore device connectivity.

Data management interface on active unit has connectivity issues

1. In Security Cloud Control , manually switch the active unit to standby. See Switch the Active Peer in the Firewall Threat Defense High Availability Pair, on page 710 .

   As an alternative, you can run the **no failover active** command on the active unit.

2. Next to the device high-availability pair you want to edit, click **Edit** ( ).
3. Choose **Routing** > **Static Route** and delete the static route defined for the old data management interface.
4. Click the **Interfaces** tab to remove the old interface configuration and configure the new data management interface.

Note Before removing the old data management interface information, remember the details if you want to use the same information.

5. Click the **High Availability** tab to update the **Monitored Interfaces** and **Standby IP Address** for the new interface.
6. Click **Save** and add the static route for the new data management interface under **Routing > Static Route** .
7. Deploy configuration changes. .
8. Re-cable the FTD when deployment reaches 90 percent so the Security Cloud Control reaches the device on the new interface.

Note After you re-cable, the deployment may fail if it timed out before re-establishing the management connection to the new interface. In that case, you must reinitiate the deployment after re-cabling for a successful deployment.

9. Verify the management connection status in Cloud-Delivered Firewall Management Center or via the **sftunnel-status-brief** command at the FTD CLI.

Internet service provider has changed

If you have changed your ISP, you can lose management connectivity, even though High Availability health is normal. Configure the new network details of the management interface using the CLI commands.

Note These commands are available only on the active unit and not on standby.

For information about the Firewall Threat Defense CLI, see the FTD command reference .

1. Connect to the device CLI and log in with the Admin username and password.
2. Use the **configure network management-data-interface** commands to update the IPv4 address, netmask, or gateway.

   Example:

```
            Configure network management-data-interface ipv4 manual 10.10.6.7
255.255.255.0 interface gig0/0
            Configuration updated successfully..!!
```

Note All other CLI commands of **configure network management-data-interface** are not supported on devices in a High Availability pair.

3. Verify the connection using **sftunnel-status-brief** .
4. In Security Cloud Control , navigate to **Security Devices > FTD**, select the device, and click **Refresh** in the **FMC Access Details** pane.
5. Click **Acknowledge** to remove the deployment block and reconcile the configuration changes.

View inventory details

The **Inventory Details** section of the **Device** page shows chassis details such as the CPU and memory.

Figure 309: Inventory details



This section shows the chassis serial number. The firewall includes two serial numbers: the chassis serial number and the PCB (circuit board) serial number. The PCB serial number is shown in the System section. The Firewall Threat Defense Virtual does not have a chassis serial number.

Note For physical Firewall Threat Defense devices, the Memory field in **Inventory Details** shows only the data plane memory, not the total system memory.

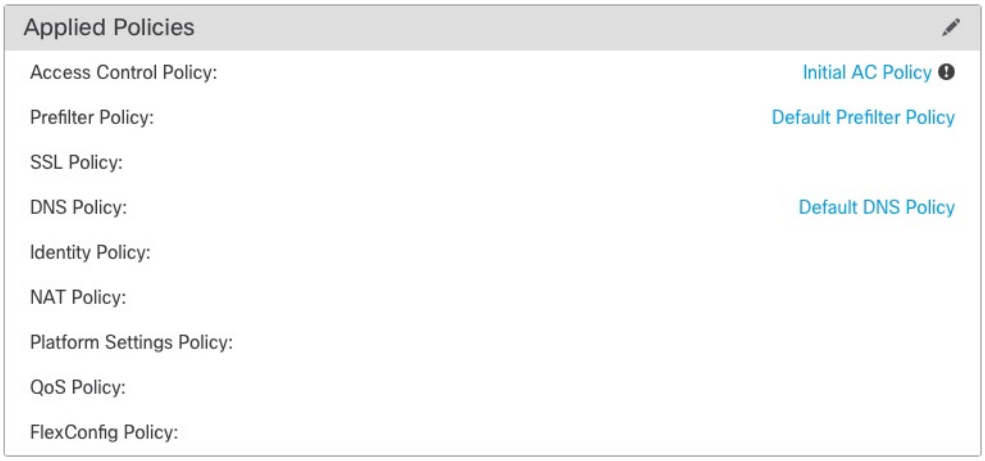
To update information, click **Refresh** ( ).

Edit applied policies

This task allows you to assign or modify policies for an individual firewall device, including access control policies and other security policies that govern device behavior.

The **Applied Policies** section of the **Device** page displays the following policies applied to your firewall:

Figure 310: Applied policies



For policies with links, you can click the link to view the policy.

For the Access Control Policy, view the **Access Policy Information for Troubleshooting** dialog box by clicking the **Exclamation** ( ) icon. This dialog box shows how access rules are expanded into access control entries (ACEs).

Figure 311: Access policy information for troubleshooting



Procedure

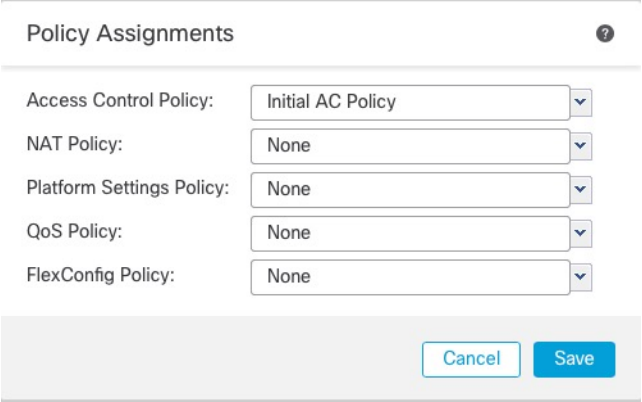
Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to assign policies, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **Applied Policies** section, click **Edit** ( ).

Figure 312: Policy assignments



Step 5 For each policy type, choose a policy from the drop-down menu. Only existing policies are listed.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes.

Edit advanced settings

The **Advanced Settings** section of the **Device** page displays a table of advanced configuration settings, as described below. You can edit any of these settings.

Table 74: Advanced settings table fields

Field	Description
Application Bypass	The state of Automatic Application Bypass on the device.
Bypass Threshold	The Automatic Application Bypass threshold, in milliseconds.

Field	Description
Object Group Search	The state of object group search on the device. While operating, the FTD device expands access control rules into multiple access control list entries based on the contents of any network or interface objects used in the access rule. You can reduce the memory required to search access control rules by enabling object group search. With object group search enabled, the system does not expand network or interface objects, but instead searches access rules for matches based on those group definitions. Object group search does not impact how your access rules are defined or how they appear in Firepower Management Center. It impacts only how the device interprets and processes them while matching connections to access control rules. **Note** By default, the **Object Group Search** is enabled when you add threat defense for the first time in the management center.
Interface Object Optimization	The state of interface object optimization on the device. During deployment, interface groups and security zones used in the access control and prefilter policies generate separate rules for each source/destination interface pair. If you enable interface object optimization, the system will instead deploy a single rule per access control/prefilter rule, which can simplify the device configuration and improve deployment performance. If you select this option, also select the **Object Group Search** option to reduce memory usage on the device.

These topics explain how to edit the advanced device settings.

Note For information about the Transfer Packets setting, see Edit general settings, on page 1133.

Configure automatic application bypass

Automatic Application Bypass (AAB) enables packet processing to continue when Snort detection becomes unavailable or exceeds processing time limits. This feature ensures network traffic flow continuity while providing diagnostic information for troubleshooting detection engine failures.

Automatic Application Bypass (AAB) allows packets to bypass detection if Snort is down or, for a Classic device, if a packet takes too long to process. AAB causes Snort to restart within ten minutes of the failure, and generates troubleshooting data that can be analyzed to investigate the cause of the Snort failure.

Caution AAB activation partially restarts the Snort process, which temporarily interrupts the inspection of a few packets.Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Note the following behavior:

Firewall Threat Defense Behavior: If Snort is down, then AAB is triggered after the specified timer duration. If Snort is up, then AAB is never triggered, even if packet processing exceeds the configured timer.

Classic Device Behavior: AAB limits the time allowed to process packets through an interface. You balance packet processing delays with your network's tolerance for packet latency.

The feature functions with any deployment; however, it is most valuable in inline deployments.

Typically, you use Rule Latency Thresholding in the intrusion policy to fast-path packets after the latency threshold value is exceeded. Rule Latency Thresholding does not shut down the engine or generate troubleshooting data.

If detection is bypassed, the device generates a health monitoring alert.

By default the AAB is disabled; to enable AAB follow these steps:

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to edit advanced device settings, click **Edit** ( ).

Step 3 Click **Device**, then click **Edit** ( ) in the **Advanced Settings** section.

Step 4 Check **Automatic Application Bypass**.

Step 5 Enter a **Bypass Threshold** from 250 ms to 60,000 ms. The default setting is 3000 milliseconds (ms).

Step 6 Click **Save**.

What to do next

- Deploy configuration changes.

Configure object group search

This task enables object group search on a Firewall Threat Defense device to reduce memory requirements for access control policies that include network or interface objects.

While operating, the Firewall Threat Defense device expands access control rules into multiple access control list entries based on the contents of any network or interface objects used in the access rule. You can reduce the memory required to search access control rules by enabling object group search. With object group search enabled, the system does not expand network or interface objects, but instead searches access rules for matches based on those group definitions. Object group search does not impact how your access rules are defined or how they appear in Cloud-Delivered Firewall Management Center. It impacts only how the device interprets and processes them while matching connections to access control rules.

Enabling object group search reduces memory requirements for access control policies that include network or interface objects. However, it is important to note that object group search might also decrease rule lookup performance and thus increase CPU utilization. You should balance the CPU impact against the reduced memory requirements for your specific access control policy. In most cases, enabling object group search provides a net operational improvement.

By default, the object group search is enabled for the threat defense devices that are added for the first time in the Cloud-Delivered Firewall Management Center. In the case of upgraded devices, if the device is configured with disabled object group search, then you need to manually enable it. You can enable it on one device at a time; you cannot enable it globally. We recommend that you enable it on a device to which you deploy access rules that use network or interface objects.



Note

- If you enable object group search and then configure and operate the device for a while, be aware that subsequently disabling the feature might lead to undesirable results.
- When you disable object group search, your existing access control rules will be expanded in the device's running configuration. If the expansion requires more memory than is available on the device, your device can be left in an inconsistent state and you might see a performance impact.
- If your device is operating normally, do not disable the object group search after you have enabled it.

Before you begin

- Model support—Firewall Threat Defense
- We recommend that you also enable transactional commit on each device. From the device CLI, enter the **asp rule-engine transactional-commit access-group** command.
- Changing this setting can be disruptive to system operation while the device recompiles the ACLs. We recommend that you change this setting during a maintenance window.
- You can use FlexConfig to configure the **object-group-search threshold** command to enable a threshold to help prevent performance degradation. When operating with a threshold, for each connection, both the source and destination IP addresses are matched against network objects. If the number of objects matched by the source address times the number matched by the destination address exceeds 10,000, the connection is dropped. Configure your rules to prevent an excessive number of matches.

Follow these steps to configure object group search:

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the Firewall Threat Defense device where you want to configure the rule, click the **Edit** (✎).

Step 3 Click the **Device** tab, then click the **Edit** (✎) in the **Advanced Settings** section.

Step 4 Check **Object Group Search**.

Step 5 To have object group search work on interface objects in addition to network objects, check **Interface Object Optimization**.

If you do not select **Interface Object Optimization**, the system deploys separate rules for each source/interface pair, rather that use the security zones and interface groups used in the rules. This means the interface groups are not available for object group search processing.

Step 6 Click **Save**.

Configure interface object optimization

Interface object optimization enables the system to deploy a single rule per access control/prefilter rule instead of generating separate rules for each source/destination interface pair, simplifying device configuration and improving deployment performance.

During deployment, interface groups and security zones used in the access control and prefilter policies generate separate rules for each source and destination interface pair. Enabling interface object optimization deploys a single rule per access control or prefilter rule. This approach simplifies device configuration and can improve deployment performance. If you select this option, also select the **Object Group Search** option to reduce evice memory usage.

Interface object optimization is disabled by default. You can enable it on one device at a time; you cannot enable it globally.

Note If you disable interface object optimization, the system deploys existing access control rules without using interface objects. This process might take longer. If object group search is enabled, its benefits do not apply to interface objects, and the device's running configuration might show expanded access control rules. If this expansion requires more memory than is available, the device can become inconsistent, which may affect performance

Before you begin

Model support—Firewall Threat Defense

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the Firewall Threat Defense device where you want to configure the rule, click the **Edit** ( ).

Step 3 Under the **Device** tab, click **Edit** ( ) in the **Advanced Settings** section.

Step 4 Check **Interface Object Optimization**.

Step 5 Click **Save**.

Edit deployment settings

Configure the deployment settings to enable auto rollback of the deployment if the management connection fails as a result of the deployment. For example, if you use data for management center access and then you misconfigure the data interface.

The **Deployment Settings** section of the **Device** page displays information for managing automatic rollback configuration. As an alternative, you can manually roll back the configuration using the **configure policy rollback** command (see Manually Roll Back the configuration if the Cloud-Delivered Firewall Management Center loses Connectivity, on page 1181).

Figure 313: Deployment settings



Table 75: Deployment settings

Field	Description
Auto Rollback Deployment if Connectivity Fails	Enabled or Disabled. You can enable auto rollback if the management connection fails as a result of the deployment; specifically if you use data for management center access, and then you misconfigure the data interface.
Connectivity Monitor Interval (in Minutes)	Shows the amount of time to wait before rolling back the configuration.

Consider the following guidelines:

- Only the previous deployment is available locally on the Firewall Threat Defense; you cannot roll back to any earlier deployments.
- Rollback is supported for high availability but not supported for clustering deployments.
- Rollback is not supported immediately after high availability creation.
- The rollback only affects configurations that you can set in the Cloud-Delivered Firewall Management Center. For example, the rollback does not affect any local configuration related to the dedicated Management interface, which you can only configure at the Firewall Threat Defense CLI. Note that if you changed data interface settings after the last Cloud-Delivered Firewall Management Center deployment using the **configure network management-data-interface** command, and then you use the rollback command, those settings will not be preserved; they will roll back to the last-deployed Cloud-Delivered Firewall Management Center settings.
- UCAPL/CC mode cannot be rolled back.
- Out-of-band SCEP certificate data that was updated during the previous deployment cannot be rolled back.
- During the rollback, connections will drop because the current configuration will be cleared.

Procedure

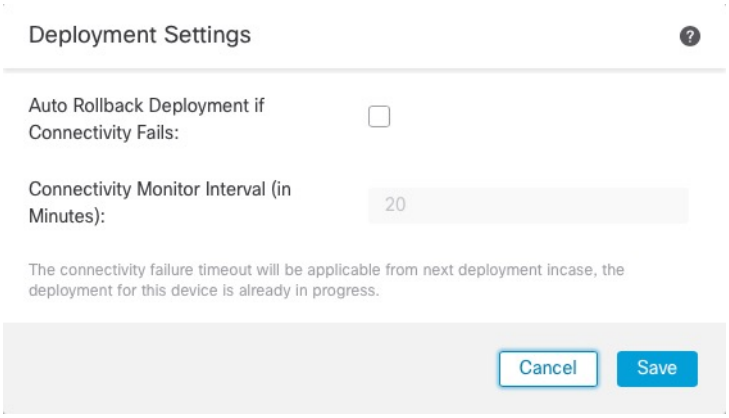
Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to assign policies, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **Deployment Settings** section, click **Edit** ( ).

Figure 314: Deployment settings



Step 5 Check **Auto Rollback Deployment if Connectivity Fails** to enable auto rollback.

Step 6 Set the **Connectivity Monitor Interval (in minutes)** to configure the wait time before rolling back the configuration. The default is 20 minutes.

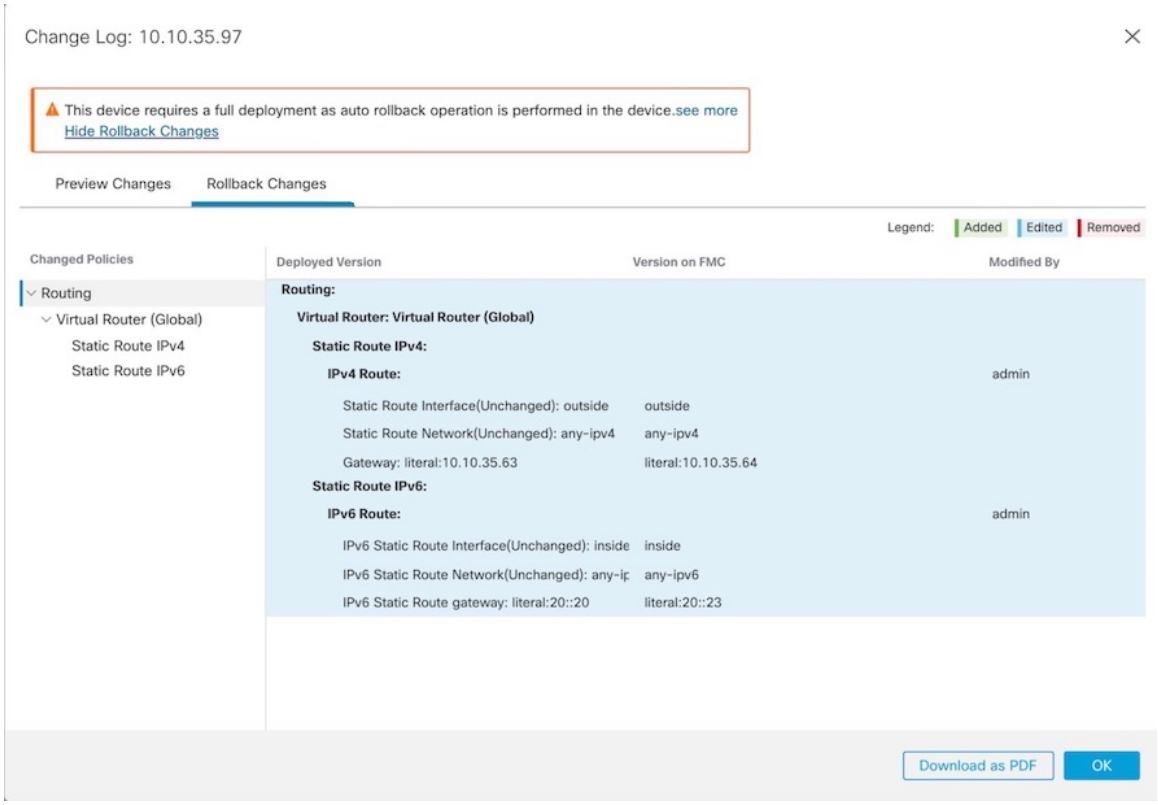
Step 7 If a rollback occurs, consider these next steps.

- If the auto rollback was successful, you see a success message instructing you to do a full deployment.
- Go to **Deploy** > **Advanced Deploy** page and click the **Preview** ( ) icon to view the parts of the configuration that were rolled back (see Deploy configuration changes, on page 202).

  To view the changes, click **Show Rollback Changes**.

  To hide the changes, click **Hide Rollback Changes**.

Figure 315: Rollback changes



- In the Deployment History Preview, you can view the rollback changes. See View deployment history, on page 210.

Step 8 Verify that the management connection was reestablished.

a. In Cloud-Delivered Firewall Management Center, check the management connection status on the connection status page.

b. Navigate to **Devices** > **Device Management** and then in **Management** area under the **Devices** tab, click **Connection Status** to view the **Connection Status** page.

c. At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

If it takes more than 10 minutes to reestablish the connection, troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 1183.

Edit cluster health monitor settings

This task allows you to modify cluster health monitor settings to control how cluster nodes monitor system health and automatically rejoin the cluster after failures.

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings for monitoring cluster node health. You can monitor any port-channel ID, any single physical interface ID, and the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link. It is always monitored.

When a topology change occurs (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature. Also, disable interface monitoring for any interfaces that are disabled. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Procedure

Step 1 Choose **Devices** > **Device Management**.

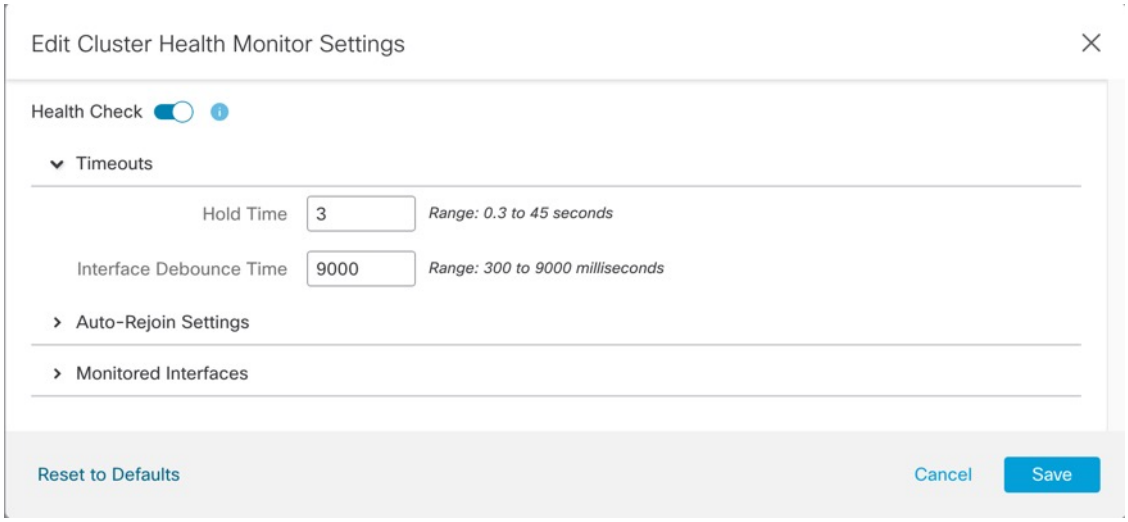
Step 2 Next to the cluster you want to modify, click **Edit** ( ).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** ( ).

Step 5 Disable the system health check by clicking the **Health Check** slider.

Figure 316: Disable the system health check



When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

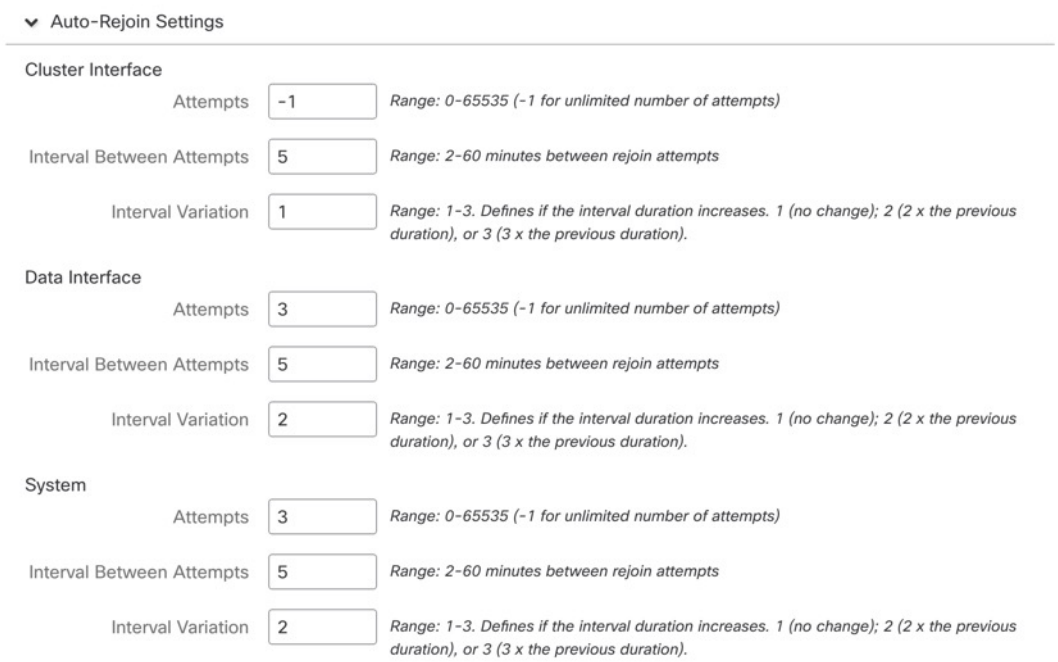
Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between 0.3 and 45 seconds. The default is 3 seconds.

- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the specified number of milliseconds before marking the interface as failed and removing the node from the cluster. If an EtherChannel transitions from a down state to an up state, such as after the switch reloads or enables an EtherChannel, a longer debounce time can prevent the interface from appearing failed on a cluster node if another cluster node bundles the ports more quickly.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 317: Configure Auto-Rejoin settings

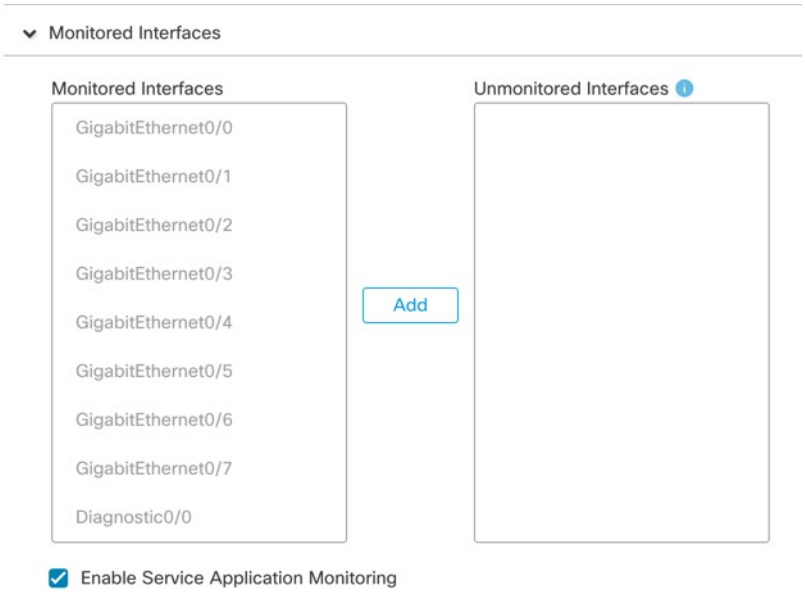


Set values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 318: Configure monitored interfaces



The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The time required to remove a member from the cluster depends on the type of interface and whether the node is established or is joining the cluster. By default, health check is enabled for all interfaces and the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Deploy configuration changes.

Hot swap an SSD

Hot swapping SSDs allows you to maintain system operation while replacing faulty drives, removing drives, or adding drives to form a RAID configuration.

If you have two SSDs, they form a RAID when you boot up. You can perform the following tasks at the Firewall Threat Defense CLI while the firewall is powered up:

- Hot swap one of the SSDs—If an SSD is faulty, you can replace it. Note that if you only have one SSD, you cannot remove it while the firewall is powered on.
- Remove one of the SSDs—If you have two SSDs, you can remove one.

- Add a second SSD—If you have one SSD, you can add a second SSD and form a RAID.

  When you use only one SSD in Secure Firewall devices, the **Drive State** will show as **degraded**.



Caution Do not remove an SSD without first removing it from the RAID using this procedure. You can cause data loss.

Procedure

Step 1 Remove one of the SSDs.

a) Remove the SSD from the RAID.

configure RAID remove-secure local-disk {**1** | **2**}

The **remove-secure** keyword removes the SSD from the RAID, disables the self-encrypting disk feature, and performs a secure erase of the SSD. If you only want to remove the SSD from the RAID and want to keep the data intact, you can use the **remove** keyword.

Example:

```
> configure raid remove-secure local-disk 2
```

b) Monitor the RAID status until the SSD no longer shows in the inventory.

show RAID

After the SSD is removed from the RAID, the **Operability** and **Drive State** will show as **degraded**. The second drive will no longer be listed as a member disk.

Example:

```
> show raid
Virtual Drive
ID:                         1
Size (MB):                  858306
Operability:                operable
Presence:                   equipped
Lifecycle:                  available
Drive State:                optimal
Type:                       raid
Level:                      raid1
Max Disks:                  2
Meta Version:               1.0
Array State:                active
Sync Action:                idle
Sync Completed:             unknown
Degraded:                   0
Sync Speed:                 none

RAID member Disk:
Device Name:                nvme0n1
Disk State:                 in-sync
Disk Slot:                  1
Read Errors:                0
```

```
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:

Device Name:                nvme1n1
Disk State:                 in-sync
Disk Slot:                  2
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:

> show raid
Virtual Drive
ID:                         1
Size (MB):                  858306
Operability:                degraded
Presence:                   equipped
Lifecycle:                  available
Drive State:                degraded
Type:                       raid
Level:                      raid1
Max Disks:                  2
Meta Version:               1.0
Array State:                active
Sync Action:                idle
Sync Completed:             unknown
Degraded:                   1
Sync Speed:                 none

RAID member Disk:
Device Name:                nvme0n1
Disk State:                 in-sync
Disk Slot:                  1
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:
```

c) Physically remove the SSD from the chassis.

Step 2 Add an SSD.

a) Physically add the SSD to the empty slot.

b) Add the SSD to the RAID.

configure RAID add local-disk {**1** | **2**}

It can take several hours to complete syncing the new SSD to the RAID, during which the firewall is completely operational. You can even reboot, and the sync will continue after it powers up. Use the **show RAID** command to show the status.

If you install an SSD that was previously used on another system, and is still locked, enter the following command:

configure RAID add local-disk {**1** | **2**} *psid*

The *psid* is printed on the label attached to the back of the SSD. Alternatively, you can reboot the system, and the SSD will be reformatted and added to the RAID.

Disable the USB port

This reference explains how to disable the USB port functionality on the device for security or operational requirements.

By default, the type-A USB port is enabled. You might want to disable USB port access for security purposes. Disabling USB is supported on these models:

- Firepower 1000 Series
- Secure Firewall 200 Series
- Secure Firewall 3100
- Secure Firewall 4200

Guidelines

- Enabling or disabling the USB port requires a reboot .
- If the USB port is disabled and you downgrade to a version that does not support this feature, the port will remain disabled, and you cannot re-enable it without erasing the NVRAM (the FXOS local-mgmt **erase secure all** command).
- If you perform a ROMMON **factory-reset** or FXOS local-mgmt **erase secure** , the USB port will be re-enabled.
- For high availability or clustering , you must disable or re-enable the port individually on each unit.

Note This feature does not affect the USB console port, if present.

Disable the USB port on a device

Disable the USB port on a device to enhance security by preventing unauthorized data transfer or malicious USB device connections.

To disable the USB port on a device, you can do so at the Firewall Threat Defense CLI.

Procedure

Step 1 Disable the USB port.

system support USB configure disable

reboot

To re-enable the USB port, enter **system support USB configure enable**.

Example:

```
>system support usb configure disable
USB Port Admin State set to 'disabled'.
Please reboot the system to apply any control state changes.

>reboot
This command will reboot the system. Continue?
Please enter 'YES' or 'NO': YES
```

Step 2 View the port status.

system support USB show

The Admin State shows the USB port configuration. The Oper State shows the current operation. For example, if you disable the USB port but do not reload, the Admin State will show disabled while the Oper State would be enabled.

Example:

```
>system support usb show
USB Port Info
---------------
Admin State: disabled
Oper State: disabled
```

Disable the USB port in multi-instance mode

This task allows you to control USB port access on your device by enabling or disabling the port through the FXOS CLI.

To disable the USB port in multi-instance mode, you can do so at the FXOS CLI.

Procedure

Step 1 Disable the USB port and reboot for the change to take effect.

a) Disable the USB port.

scope fabric-interconnect

disable USB-port

commit buffer

b) Reboot the chassis.

connect local-mgmt

reboot

Example:

```
firepower-4245 /fabric-interconnect # disable usb-port
Note: USB enablement or disablement changes are effected only after FXOS reboot.
Confirm change? (yes/no) [yes]:
device /fabric-interconnect* # commit buffer
```

```
Note: USB enablement or disablement changes are effected only after FXOS reboot.
Confirm change? (yes/no) [yes]:yes
firepower-4245 /fabric-interconnect # connect local-mgmt
firepower-4245(local-mgmt)# reboot
Before rebooting, please take a configuration backup.
Do you still want to reboot? (yes/no):yes
Broadcast message from admin@firepower-4245 (Wed Feb 21 05:59:55 2024):
All shells being terminated due to system /sbin/reboot
```

Step 2 Enable the USB port and reboot for the change to take effect.

a) Enable the USB port.

scope fabric-interconnect

enable USB-port

commit buffer

b) Reboot the chassis.

connect local-mgmt

reboot

Example:

```
firepower-4245 /fabric-interconnect # enable usb-port
Note: USB enablement or disablement changes are effected only after FXOS reboot.
Confirm change? (yes/no) [yes]:
device /fabric-interconnect* # commit buffer
Note: USB enablement or disablement changes are effected only after FXOS reboot.
Confirm change? (yes/no) [yes]:yes
firepower-4245 /fabric-interconnect # connect local-mgmt
firepower-4245(local-mgmt)# reboot
Before rebooting, please take a configuration backup.
Do you still want to reboot? (yes/no):yes
Broadcast message from admin@firepower-4245 (Wed Feb 21 05:59:55 2024):
All shells being terminated due to system /sbin/reboot
```

Step 3 View the USB port status.

scope fabric-interconnect

show USB-port

The Admin State shows the USB port configuration. The Oper State shows the current operation. For example, if you disable the USB port but do not reload, the Admin State will show Disabled while the Oper State would will Enabled.

Example:

```
firepower-4245# scope fabric-interconnect
firepower-4245 /fabric-interconnect # show usb-port
Usb Port:
Equipment         Admin State Oper State
----------------- ----------- ----------
A                 Disabled    Disabled
```

Configure SNMP for FXOS

You can configure the Simple Network Management Protocol (SNMP) for the underlying FXOS operating system that controls low-level chassis functions. For SNMP for the Firewall Threat Defense, see Configure SNMP for a security device, on page 1264.

SNMP is an application-layer protocol that provides a message format for communication between SNMP managers and agents. SNMP provides a standardized framework and a common language used for the monitoring and management of devices in a network.

The SNMP framework consists of three parts:

- An SNMP manager—The system used to control and monitor the activities of network devices using SNMP.
- An SNMP agent—The software component within the chassis that maintains the data for the chassis and reports the data, as needed, to the SNMP manager. The chassis includes the agent and a collection of MIBs. To enable the SNMP agent and create the relationship between the manager and agent, enable and configure SNMP in the Cloud-Delivered Firewall Management Center.
- A managed information base (MIB)—The collection of managed objects on the SNMP agent.

The chassis supports SNMPv1, SNMPv2c and SNMPv3. Both SNMPv1 and SNMPv2c use a community-based form of security. SNMPv3 uses username and password for security.

Enable SNMP for FXOS

Enable SNMP and configure settings.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **SNMP**.

Step 3 Complete the following fields:

Name	Description
Admin State check box	Whether SNMP is enabled or disabled. Enable this service only if your system includes integration with an SNMP server.
Port field	The port on which the Firepower chassis communicates with the SNMP host. You cannot change the default port.

Name	Description
Community field	The default SNMPv1 or v2 community name or SNMP v3 username and password that the Firepower chassis includes on any trap messages it sends to the SNMP host. Enter a valid community string for SNMPv1 and SNMPv2: • Alphanumeric string between 1 and 32 characters and special characters ! (exclamation), - (hyphen), ~ (tilde), && (double ampersand), [ ] (square brackets), ^ (carat), ' (single quote), " (double quotes), and < > (angle brackets). • Do not use @ (at sign), \ (backslash), ? (question mark) or an empty space. • The string can also be in ASCII characters ranging 0x21 to 0x7E inclusive, excluding HTML interjection vectors, namely single quote ('), double quotes ("), and angle brackets (< >). Enter a valid username and password for SNMPv3: • Username can be alphanumeric string and can include @ (at sign), \ (backslash), . (period), _ (underscore), and - (hyphen). • The password restrictions are same as the community string restrictions. Note that if the **Community** field is already set, the text to the right of the empty field reads **Set: Yes**. If the **Community** field is not yet populated with a value, the text to the right of the empty field reads **Set: No**.
System Admin Name field	The contact person responsible for the SNMP implementation. Enter a string of up to 255 characters, such as an email address or a name and telephone number.
Location field	The location of the host on which the SNMP agent (server) runs. Enter an alphanumeric string up to 510 characters.

Step 4 Click **Save**.

What to do next

Create SNMP traps and users.

Create an SNMP Trap for FXOS

Create traps for SNMP.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **SNMP**.

Step 3 In the **SNMP Traps Configuration** area, click **Add**.

Step 4 In the **SNMP Trap Configuration** dialog box, complete the following fields:

Name	Description
Host Name field	The hostname or IP address of the SNMP host to which the Firepower chassis should send the trap.
Community field	The SNMP v1 or v2 community name or the SNMP v3 username the Firepower chassis includes when it sends the trap to the SNMP host. This must be the same as the community or username that is configured for the SNMP service. Enter an alphanumeric string between 1 and 32 characters. Do not use @ (at sign), \ (backslash), " (double quote), ? (question mark) or an empty space.
Port field	The port on which the Firepower chassis communicates with the SNMP host for the trap. Enter an integer between 1 and 65535.
Version field	The SNMP version and model used for the trap. This can be one of the following: • **V1** • **V2** • **V3**
Type field	If you select **V2** or **V3** for the version, the type of trap to send. This can be one of the following: • **Traps** • **Informs**
Privilege field	If you select **V3** for the version, the privilege associated with the trap. This can be one of the following: • **Auth**—Authentication but no encryption • **Noauth**—No authentication or encryption • **Priv**—Authentication and encryption

Step 5 Click **OK** to close the **SNMP Trap Configuration** dialog box.

Step 6 Click **Save**.

Create an SNMP user for FXOS

Create SNMP users.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **SNMP**.

Step 3 In the **SNMP Users Configuration** area, click **Add**.

Step 4 In the **SNMP User Configuration** dialog box, complete the following fields:

Name	Description
Username field	The username assigned to the SNMP user. Enter up to 32 letters or numbers. The name must begin with a letter and you can also specify _ (underscore), . (period), @ (at sign), and - (hyphen).
Auth Algorithm Type field	The authorization type: **SHA**.
Use AES-128 checkbox	If checked, this user uses AES-128 encryption. **Note** SNMPv3 does not support DES. If you leave the AES-128 box unchecked, no privacy encryption will be done and any configured privacy password will have no effect.
Authentication Password field	The password for the user.
Confirm field	The password again for confirmation purposes.
Encryption Password field	The privacy password for the user.
Confirm field	The privacy password again for confirmation purposes.

Step 5 Click **OK** to close the **SNMP User Configuration** dialog box.

Step 6 Click **Save**.

Configure alarms for the ISA 3000

You can configure the alarm system on a Cisco ISA 3000 device to alert you when undesirable conditions occur.

About alarms

You can configure the ISA 3000 to issue alarms for a variety of conditions. If any conditions do not match the configured settings, the system triggers an alarm, which is reported by way of LEDs, syslog messages, SNMP traps, and through external devices connected to the alarm output interface. By default, triggered alarms issue syslog messages only.

You can configure the alarm system to monitor the following:

- Power supply.
- Primary and secondary temperature sensors.
- Alarm input interfaces.

The ISA 3000 has internal sensors plus two alarm input interfaces and one alarm output interface. You can connect external sensors, such as door sensors, to the alarm inputs. You can connect external alarm devices, such as buzzers or lights, to the alarm output interface.

The alarm output interface is a relay mechanism. Depending on the alarm conditions, the relay is either energized or de-energized. When it is energized, any device connected to the interface is activated. A de-energized relay results in the inactive state of any connected devices. The relay remains in an energized state as long as alarms are triggered.

For information about connecting external sensors and the alarm relay, Refer to Alarms for the Cisco ISA 3000 chapter in Cisco ASA Series General Operations CLI Configuration Guide. see Cisco ISA 3000 Industrial Security Appliance Hardware Installation Guide.

Alarm input interfaces

You can connect the alarm input interfaces (or contacts) to external sensors, such as one that detects if a door is open.

Each alarm input interface has a corresponding LED. These LEDs convey the alarm status of each alarm input. You can configure the trigger and severity for each alarm input. In addition to the LED, you can configure the contact to trigger the output relay (to activate an external alarm), to send syslog messages, and to send SNMP traps.

The following table explains the statuses of the LEDs in response to alarm conditions for the alarm inputs. It also explains the behavior for the output relay, syslog messages, and SNMP traps, if you enable these responses to the alarm input.

Alarm Status	LED	Output Relay	Syslog	SNMP Trap
Alarm not configured	Off	—	—	—
No alarms triggered	Solid green	—	—	—
Alarm activated	Minor alarm—solid red Major alarm—flashing red	Relay energized	Syslog generated	SNMP trap sent
Alarm end	Solid green	Relay de-energized	Syslog generated	—

Alarm output interface

You can connect an external alarm, such as a buzzer or light, to the alarm output interface.

The alarm output interface functions as a relay and also has a corresponding LED, which conveys the alarm status of an external sensor connected to the input interface, and internal sensors such as the dual power supply and temperature sensors. You configure which alarms should activate the output relay, if any.

The following table explains the statuses of the LEDs and output relay in response to alarm conditions. It also explains the behavior for syslog messages, and SNMP traps, if you enable these responses to the alarm.

Alarm Status	LED	Output Relay	Syslog	SNMP Trap
Alarm not configured	Off	—	—	—
No alarms triggered	Solid green	—	—	—
Alarm activated	Solid red	Relay energized	Syslog generated	SNMP trap sent
Alarm end	Solid green	Relay de-energized	Syslog generated	—

Syslog alarms

By default, the system sends syslog messages when any alarm is triggered. You can disable syslog messaging if you do not want the messages.

For syslog alarms to work, you must also enable diagnostic logging. Choose **Devices** > **Platform Settings**, add or edit a Threat Defense platform settings policy that is assigned to the device, and configure destinations and settings on the **Syslog** page. For example, you can configure a syslog server, console logging, or internal buffer logging.

Without enabling a destination for diagnostic logging, the alarm system has nowhere to send syslog messages.

SNMP Alarms

You can optionally configure the alarms to send SNMP traps to your SNMP server. For SNMP trap alarms to work, you must also configure SNMP settings.

Choose **Devices** > **Platform Settings**, add or edit a Threat Defense platform settings policy that is assigned to the device, and enable SNMP and configure settings on the **SNMP** page.

Defaults for alarms

The following table specifies the defaults for alarm input interfaces (contacts), redundant power supply, and temperature.

	Alarm	Trigger	Severity	SNMP Trap	Output Relay	Syslog Message
Alarm Contact 1	Enabled	Closed State	Minor	Disabled	Disabled	Enabled
Alarm Contact 2	Enabled	Closed State	Minor	Disabled	Disabled	Enabled

	Alarm	Trigger	Severity	SNMP Trap	Output Relay	Syslog Message
Redundant Power Supply (when enabled)	Enabled	—	—	Disabled	Disabled	Enabled
Temperature	Enabled for the primary temperature alarm (default values of 92°C and -40°C for the high and low thresholds respectively) Disabled for the secondary alarm.	—	—	Enabled for primary temperature alarm	Enabled for primary temperature alarm	Enabled for primary temperature alarm

Prerequisites for alarms

Model support

Firewall Threat Defense on the ISA 3000.

Configure alarms for the ISA 3000

You use FlexConfig to configure alarms for the ISA 3000. The following topics explain how to configure the different types of alarms.

Configure alarm input contacts

If you connect the alarm input contacts (interfaces) to external sensors, you can configure the contacts to issue alarms based on the input from the sensor. In fact, the contacts are enabled by default to send syslog messages if the contact is closed, that is, if the electrical current stops flowing through the contact. You need to configure the contact only if the defaults do not meet your requirements.

The alarm contacts are numbered 1 and 2, so you need to understand how you have wired the physical pins to configure the correct settings. You configure the contacts separately.

Procedure

Step 1 Create the FlexConfig object to configure the alarm input contacts.

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Configure_Alarm_Contacts.

- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features.
- **Object body**—In the object body, type the commands needed to configure the alarm contacts. The following steps explain the commands.

c) Configure a description for the alarm contact.

alarm contact {**1** | **2**} **description** *string*

For example, to set the description of contact 1 to "Door Open," enter the following:

```
alarm contact 1 description Door Open
```

d) Configure the severity for the alarm contact.

alarm contact {**1** | **2** | **any**} **severity** {**major** | **minor** | **none**}

Instead of configuring one contact, you can specify **any** to change the severity for all contacts. The severity controls the behavior of the LED associated with the contact.

- **major**—The LED blinks red.
- **minor**—The LED is solid red. This is the default.
- **none**—The LED is off.

For example, to set the severity of contact 1 to Major, enter the following:

```
alarm contact 1 severity major
```

e) Configure the trigger for the alarm contact.

alarm contact {**1** | **2** | **any**} **trigger** {**open** | **closed**}

Instead of configuring one contact, you can specify **any** to change the trigger for all contacts. The trigger determines the electrical condition that signals an alert.

- **open**—The normal condition for the contact is closed, that is, the electrical current is running through the contact. An alert is triggered if the contact becomes open, that is, the electrical current stops flowing.
- **closed**—The normal condition for the contact is open, that is, the electrical current does not run through the contact. An alert is triggered if the contact becomes closed, that is, the electrical current starts running through the contact. This is the default.

For example, you connect a door sensor to alarm input contact 1, and its normal state has no electrical current flowing through the alarm contact (it is open). If the door is opened, the contact is closed and electrical current flows through the alarm contact. You would set the alarm trigger to closed so that the alarm goes off when the electrical current starts flowing.

```
alarm contact 1 trigger closed
```

f) Configure the actions to take when the alarm contact is triggered.

alarm facility input-alarm {**1** | **2**} {**relay** | **syslog** | **notifies**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message. This option is enabled by default.
- **notifies**—Send an SNMP trap.

For example, to enable all actions for the alarm input contact 1, enter the following:

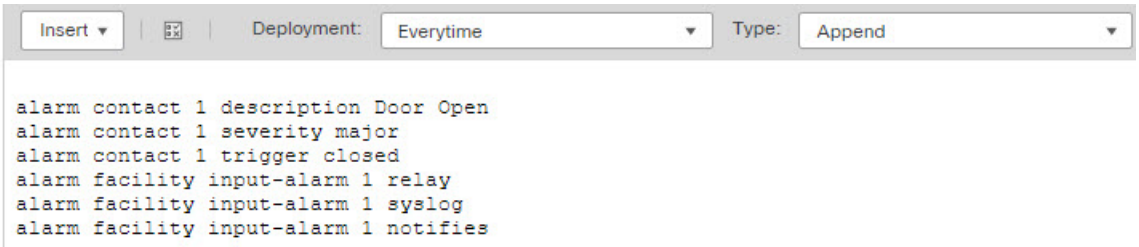
```
alarm facility input-alarm 1 relay
alarm facility input-alarm 1 syslog
alarm facility input-alarm 1 notifies
```

g) Verify that the object body contains the commands you want.

For example, if your template includes all of the command examples shown in this procedure, the object body would have the following commands:

```
alarm contact 1 description Door Open
alarm contact 1 severity major
alarm contact 1 trigger closed
alarm facility input-alarm 1 relay
alarm facility input-alarm 1 syslog
alarm facility input-alarm 1 notifies
```

The object body should look similar to the following:



h) Click **Save**.

Step 2 Create the FlexConfig policy and assign it to the devices.

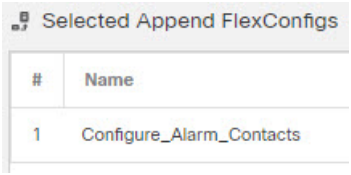
a) Choose **Devices** > **FlexConfig**.

b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the alarm contact FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.



d) Click **Save**.

e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the FlexConfig object look correct. These will be shown at the end of the preview. Note that you will also see commands generated from other changes you have made to managed features. For the alarm contact commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
alarm contact 1 description Door Open
alarm contact 1 severity major
alarm contact 1 trigger closed
alarm facility input-alarm 1 relay
alarm facility input-alarm 1 syslog
alarm facility input-alarm 1 notifies
```

Step 3 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the deployed configuration, on page 1342.

Configure power supply alarms

The ISA 3000 has two power supplies. By default, the system operates in single-power mode. However, you can configure the system to operate in dual mode, where the second power supply automatically provides power if the primary power supply fails. When you enable dual-mode, the power supply alarm is automatically enabled to send syslog alerts, but you can disable the alert altogether, or also enable SNMP traps or the alarm hardware relay.

The following procedure explains how to enable dual mode, and how to configure the power supply alarms.

Procedure

Step 1 Create the FlexConfig object to configure the power supply alarm.

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Power_Supply_Alarms.

- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features.
- **Object body**—In the object body, type the commands needed to configure the power supply alarms. The following steps explain the commands.

c) Enable dual power supply mode.

power-supply dual

For example:

```
power-supply dual
```

d) Configure the actions to take when the power supply alarm is triggered.

alarm facility power-supply rps {**relay** | **syslog** | **notifies** | **disable**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message. This option is enabled by default.
- **notifies**—Send an SNMP trap.
- **disable**—Disable the power supply alarm. Any other actions configured for the power supply alarm are inoperable.

For example, to enable all actions for the power supply alarm, enter the following:

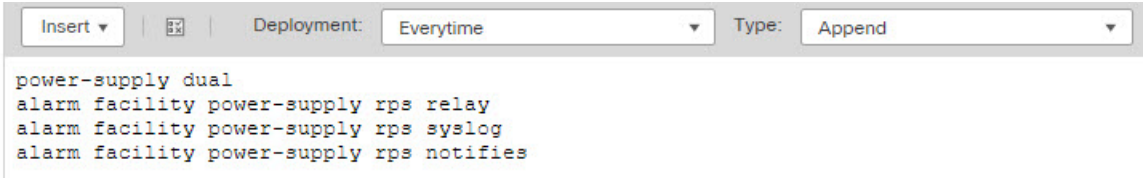
```
alarm facility power-supply rps relay
alarm facility power-supply rps syslog
alarm facility power-supply rps notifies
```

e) Verify that the object body contains the commands you want.

For example, if your template includes all of the command examples shown in this procedure, the object body would have the following commands:

```
power-supply dual
alarm facility power-supply rps relay
alarm facility power-supply rps syslog
alarm facility power-supply rps notifies
```

The object body should look similar to the following:



f) Click **Save**.

Step 2 Create the FlexConfig policy and assign it to the devices.

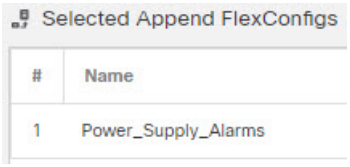
a) Choose **Devices** > **FlexConfig**.

b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the power supply alarm FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.



d) Click **Save**.

e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the FlexConfig object look correct. These will be shown at the end of the preview. Note that you will also see commands generated from other changes you have made to managed features. For the power supply alarm commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
power-supply dual
alarm facility power-supply rps relay
alarm facility power-supply rps syslog
alarm facility power-supply rps notifies
```

Step 3 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the deployed configuration, on page 1342.

Configure temperature alarms

You can configure alarms based on the temperature of the CPU card in the device.

You can set a primary and secondary temperature range. If the temperature drops below the low threshold, or exceeds the high threshold, the alarm is triggered.

The primary temperature alarm is enabled by default for all alarm actions: output relay, syslog, and SNMP. The default settings for the primary temperature range is -40°C to 92°C.

The secondary temperature alarm is disabled by default. You can set the secondary temperature within the range -35°C to 85°C.

Because the secondary temperature range is more restrictive than the primary range, if you set either the secondary low or high temperature, that setting disables the corresponding primary setting, even if you configure non-default values for the primary setting. You cannot enable two separate high and two separate low temperature alarms.

Thus, in practice, you should configure the primary only, or the secondary only, setting for high and low.

Procedure

Step 1 Create the FlexConfig object to configure the temperature alarms.

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Configure_Temperature_Alarms.
- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features.
- **Object body**—In the object body, type the commands needed to configure the temperature alarms. The following steps explain the commands.

c) Configure the acceptable temperature range.

alarm facility temperature {**primary** | **secondary**} {**low** | **high**} *temperature*

The temperature is in Celsius. The allowed range for the primary alarm is -40 to 92, which is also the default range. The allowed range for the secondary alarm is -35 to 85. The low value must be lower than the high value.

For example, to set a more restrictive temperature range of -20 to 80, which falls within the allowed range for the secondary alarm, configure the secondary alarm as follows:

```
alarm facility temperature secondary low -20
alarm facility temperature secondary high 80
```

d) Configure the actions to take when the temperature alarm is triggered.

alarm facility temperature {**primary** | **secondary**} {**relay** | **syslog** | **notifies**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message.

- **notifies**—Send an SNMP trap.

For example, to enable all actions for the secondary temperature alarm, enter the following:

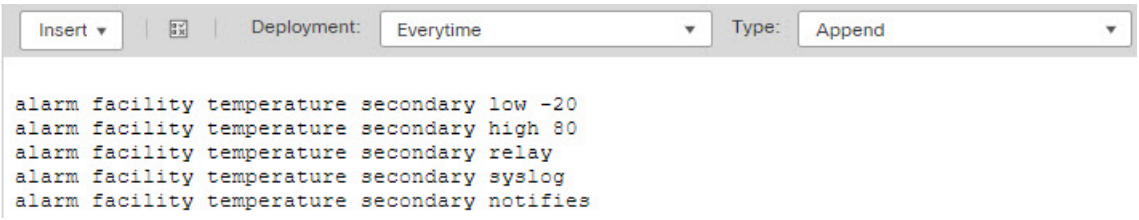
```
alarm facility temperature secondary relay
alarm facility temperature secondary syslog
alarm facility temperature secondary notifies
```

e) Verify that the object body contains the commands you want.

For example, if your template includes all of the command examples shown in this procedure, the object body would have the following commands:

```
alarm facility temperature secondary low -20
alarm facility temperature secondary high 80
alarm facility temperature secondary relay
alarm facility temperature secondary syslog
alarm facility temperature secondary notifies
```

The object body should look similar to the following:



f) Click **Save**.

Step 2 Create the FlexConfig policy and assign it to the devices.

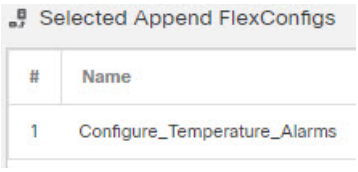
a) Choose **Devices** > **FlexConfig**.

b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the temperature alarms FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.



d) Click **Save**.

e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the FlexConfig object look correct. These will be shown at the end of the

preview. Note that you will also see commands generated from other changes you have made to managed features. For the temperature alarms commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
alarm facility temperature secondary low -20
alarm facility temperature secondary high 80
alarm facility temperature secondary relay
alarm facility temperature secondary syslog
alarm facility temperature secondary notifies
```

Step 3 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the deployed configuration, on page 1342.

Turn off the external alarm

If you are using an external alarm that is attached to the alarm output, and the alarm is triggered, you can turn off the external alarm from the device CLI using the **clear facility-alarm output** command. This command de-energizes the output pin and also turns off the output LED.

Monitoring alarms

The following topics explain how to monitor and manage alarms.

Monitoring alarm status

You can use the following commands in the CLI to monitor alarms.

- **show alarm settings**

  Shows the current configuration for each possible alarm.

- **show environment alarm-contact**

  Shows information about the physical status of the input alarm contacts.

- **show facility-alarm relay**

  Shows information about the alarms that have triggered the output relay.

- **show facility-alarm status** [**info** | **major** | **minor**]

  Shows information on all alarms that have been triggered. You can limit the view by filtering on **major** or **minor** status. The **info** keyword provides the same output as using no keyword.

Monitoring syslog messages for alarms

Depending on the type of alarms you configure, you might see the following syslog messages.

Dual Power Supply Alarms

- %FTD-1-735005: Power Supply Unit Redundancy OK
- %FTD-1-735006: Power Supply Unit Redundancy Lost

Temperature Alarms

In these alarms, *Celsius* is replaced by the temperature detected on the device, in Celsius.

- %FTD-6-806001: Primary alarm CPU temperature is High *Celsius*
- %FTD-6-806002: Primary alarm for CPU high temperature is cleared
- %FTD-6-806003: Primary alarm CPU temperature is Low *Celsius*
- %FTD-6-806004: Primary alarm for CPU Low temperature is cleared
- %FTD-6-806005: Secondary alarm CPU temperature is High *Celsius*
- %FTD-6-806006: Secondary alarm for CPU high temperature is cleared
- %FTD-6-806007: Secondary alarm CPU temperature is Low *Celsius*
- %FTD-6-806008: Secondary alarm for CPU Low temperature is cleared

Alarm Input Contact Alarms

In these alarms, *description* is the description for the contact that you configured.

- %FTD-6-806009: Alarm asserted for ALARM_IN_1 *alarm_1_description*
- %FTD-6-806010: Alarm cleared for ALARM_IN_1 *alarm_1_description*
- %FTD-6-806011: Alarm asserted for ALARM_IN_2 *alarm_2_description*
- %FTD-6-806012: Alarm cleared for ALARM_IN_2 *alarm_2_description*

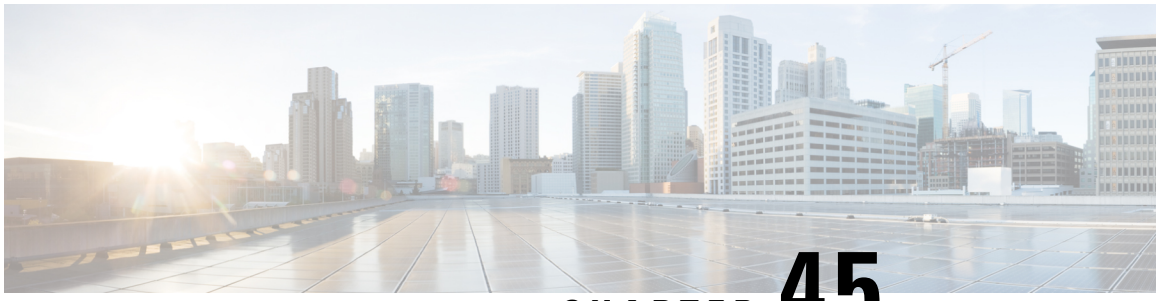
History for Device Settings

Provides a comprehensive reference of feature introductions, command modifications, and interface updates related to device settings and management.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Recovery-config mode for emergency on-device configuration and out-of-band configuration detection on the Cloud-Delivered Firewall Management Center	20250219	7.7.0	If you lose the management connection to your device, you can make select configuration changes directly at the device CLI to: • Restore the management connection if you are using a data interface for manager access • Make select policy changes that can't wait until the connection is restored After the management connection is restored, the Cloud-Delivered Firewall Management Center will detect the configuration changes on the device. It does not automatically update the device configuration in the Cloud-Delivered Firewall Management Center ; you must view the configuration differences, acknowledge that the device configuration is different, and then manually make the same changes in the Cloud-Delivered Firewall Management Center before you deploy. New/modified diagnostic CLI (**system support diagnostic-cli**) command: **configure recovery-config** New/modified screens: **Devices** > **Device Management** , click **Edit** ( ) , then **Device** > **Health** > **Out of Band Status**
High availability is supported with redundant manager access data interfaces	20250219	7.7.0	You can now use redundant manager access data interfaces with high availability.
Cluster health monitor settings.	20221213	Any	You can now edit cluster health monitor settings. New/modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Redundant manager access data interface.	20221213	7.3.0	When you use a data interface for manager access, you can configure a secondary data interface to take over management functions if the primary interface goes down. The device uses SLA monitoring to track the viability of the static routes and an ECMP zone that contains both interfaces so management traffic can use both interfaces. New/modified screens: • **Devices** > **Device Management** > **Device** > **Management** • **Devices** > **Device Management** > **Device** > **Interfaces** > **Manager Access**
Policy rollback support for high availability devices.	20220609	7.2.0	The **configure policy rollback** command is supported for high availability devices.
Auto rollback of a deployment that causes a loss of management connectivity.	20220609		You can now enable auto rollback of the configuration if a deployment causes the management connection between the management center and the threat defense to go down. Previously, you could only manually rollback a configuration using the **configure policy rollback** command. New/modified screens: • **Devices** > **Device Management** > **Device** > **Deployment Settings** • **Deploy** > **Advanced Deploy** > **Preview** • **Deploy** > **Deployment History** > **Preview**
Object group search is enabled by default for access control rules.	20220609	7.2.0	The **Object Group Search** setting is enabled by default for managed devices starting with Version 7.2.0. This option is in the **Advanced Settings** section when editing device settings on the Device Management page.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Import and export device configurations.	20220609	7.2.0	You can export the device-specific configuration, and you can then import the saved configuration for the same device in the following use cases: • Moving the device to a different FMC. • Restore an old configuration. • Reregistering a device. New/modified screens: **Devices** > **Device Management** > **Device** > **General**
Update the FMC IP address on FTD.	20220609	7.0.3	If you change the FMC IP address, you can now use the FTD CLI to update the device. New/modified commands: **configure manager edit**
Alarms for the Cisco ISA 3000 series.	6.7	Any	Configuring alarms for the Cisco ISA 3000 series was validated using FlexConfig. You should be able to configure the alarms in older releases that support FlexConfig, except for the dual power supply alarms. Supported platforms: Secure Firewall Threat Defense



CHAPTER 45

DHCP and DDNS

The following topics explain DHCP and DDNS services and how to configure them on Threat Defense devices.

- About DHCP and DDNS services, on page 1223
- Requirements and prerequisites for DHCP and DDNS, on page 1225
- Guidelines for DHCP and DDNS services, on page 1225
- Configure the DHCPv4 server, on page 1227
- DHCPv6 stateless server configuration, on page 1229
- Configure the DHCP relay agent, on page 1233
- Configure Dynamic DNS, on page 1234
- History for DHCP and DDNS, on page 1239

About DHCP and DDNS services

This topic describes the DHCP server, DHCP relay agent, and DDNS update.

DHCPv4 servers

A DHCPv4 server provides network configuration parameters, such as IP addresses, to DHCP clients.

- delivers configuration parameters directly to DHCP clients attached to Firewall Threat Defense device interfaces, and
- uses UDP port communication for client-server messaging.

DHCPv4 server communication details

The DHCP communication uses specific network protocols and port assignments:

- IPv4 DHCP client: Uses broadcast rather than multicast address to reach the server and listens for messages on UDP port 68
- DHCP server: Listens for messages on UDP port 67

DHCP options

DHCP provides a framework for passing configuration information to hosts on a TCP/IP network. The configuration parameters are carried in tagged items that are stored in the Options field of the DHCP message and the data are also called options. Vendor information is also stored in Options, and all of the vendor information extensions can be used as DHCP options.

Common DHCP options

Cisco IP Phones download their configuration from a TFTP server. When a Cisco IP Phone starts, if it does not have both the IP address and TFTP server IP address preconfigured, it sends a request with option 150 or 66 to the DHCP server to obtain this information.

- DHCP option 150 provides the IP addresses of a list of TFTP servers.
- DHCP option 66 gives the IP address or the hostname of a single TFTP server.
- DHCP option 3 sets the default route.

A single request might include both options 150 and 66. In this case, the DHCP server provides values for both options in the response if they are already configured on the Firewall Threat Defense.

You can use advanced DHCP options to provide DNS, WINS, and domain name parameters to DHCP clients. DHCP option 15 is used for the DNS domain suffix. You can also use the DHCP automatic configuration setting to obtain these values or define them manually. When you use more than one method to define this information, it is passed to DHCP clients in the following sequence:

1. Manually configured settings.
2. Advanced DHCP options settings.
3. DHCP automatic configuration settings.

Domain name configuration precedence

You can manually define the domain name that you want the DHCP clients to receive and then enable DHCP automatic configuration. Although DHCP automatic configuration discovers the domain together with the DNS and WINS servers, the manually defined domain name is passed to DHCP clients with the discovered DNS and WINS server names, because the domain name discovered by the DHCP automatic configuration process is superseded by the manually defined domain name.

DHCPv6 stateless server

A DHCPv6 stateless server is a DHCP IPv6 configuration that

- provides information such as DNS server or domain name when clients send Information Request (IR) packets,
- accepts IR packets but does not assign addresses to clients, and
- works in conjunction with StateLess Address Auto Configuration (SLAAC) and Prefix Delegation features.

DHCPv6 stateless server operation

For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature (Enable the IPv6 prefix delegation client, on page 1093), you can configure the Firewall Threat Defense to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the Firewall Threat Defense by defining a DHCP IPv6 Pool and assigning it to the DHCPv6 server. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses based on prefixes received in Router Advertisement messages; in other words, based on the prefix that the Firewall Threat Defense received using Prefix Delegation.

DHCP relay agents

A DHCP relay agent forwards DHCP requests received on an interface to one or more DHCP servers.

DHCP relay agent operation

DHCP clients use UDP broadcasts to send their initial DHCPDISCOVER messages because they do not have information about the network to which they are attached. If the client is on a network segment that does not include a server, UDP broadcasts normally are not forwarded by the Firewall Threat Defense device because it does not forward broadcast traffic. The DHCP relay agent lets you configure the interface of the Firewall Threat Defense device that is receiving the broadcasts to forward DHCP requests to a DHCP server on another device.

Requirements and prerequisites for DHCP and DDNS

This reference provides the model support and user role requirements necessary for configuring DHCP and DDNS functionality.

Model support

Firewall Threat Defense

User roles

- Admin
- Access Admin
- Network Admin

Guidelines for DHCP and DDNS services

Context mode restrictions

DHCPv6 stateless server is not supported in multiple context mode.

Firewall mode restrictions

Consider these firewall mode limitations when configuring DHCP and DDNS services:

- DHCP Relay is not supported in transparent firewall mode or in routed mode on the BVI or bridge group member interface.
- DHCP Server is supported in transparent firewall mode on a bridge group member interface. In routed mode, the DHCP server is supported on the BVI interface, not the bridge group member interface. The BVI must have a name for the DHCP server to operate.
- DDNS is not supported in transparent firewall mode or in routed mode on the BVI or bridge group member interface.
- DHCPv6 stateless server is not supported in transparent firewall mode or in routed mode on the BVI or bridge group member interface.

Clustering limitations

DHCPv6 stateless server is not supported with clustering.

IPv6 support

DHCPv4 server guidelines

Follow these guidelines when configuring DHCPv4 server:

- The maximum available DHCP pool is 256 addresses.
- You can configure a DHCP server on any interface with a name and IP address, such as a physical interface, a subinterface, or a BVI in routed mode.
- Do not select an interface that uses DHCP or PPPoE to obtain it's address as a DHCP server interface.
- You can configure only one DHCP server on each interface. Each interface can have its own pool of addresses to use. However the other DHCP settings, such as DNS servers, domain name, options, ping timeout, and WINS servers, are configured globally and used by the DHCP server on all interfaces.
- You cannot configure an interface as a DHCP client if that interface also has DHCP server enabled; you must use a static IP address.
- You cannot configure both a DHCP server and DHCP relay on the same device, even if you want to enable them on different interfaces; you can only configure one type of service.
- The Firewall Threat Defense does not support DHCP clients behind DHCP relay servers; the client must be directly connected to the Firewall Threat Defense.
- The DHCP server does not support BOOTP requests.

DHCPv6 server guidelines

Consider these guidelines when configuring DHCPv6 server:

- The DHCPv6 Stateless server cannot be configured on an interface where the DHCPv6 address, Prefix Delegation client, or DHCPv6 relay is configured.

DHCP relay guidelines

Follow these guidelines when configuring DHCP relay:

- You can configure a maximum of 10 DHCPv4 relay servers per virtual router, global (VRF) and interface-specific servers combined, with a maximum of 4 servers per interface.
- You can configure a maximum of 10 DHCPv6 relay servers per virtual router. Interface-specific servers for IPv6 are not supported.
- You cannot configure both a DHCP server and DHCP relay on the same device, even if you want to enable them on different interfaces; you can only configure one type of service.
- Do not select an interface that uses DHCP or PPPoE to obtain it's address as a DHCP relay interface.
- DHCP relay services are not available in transparent firewall mode or in routed mode on the BVI or bridge group member interface. You can, however, allow DHCP traffic through using an access rule. To allow DHCP requests and replies through the Firewall Threat Defense, you need to configure two access rules, one that allows DCHP requests from the inside interface to the outside (UDP destination port 67), and one that allows the replies from the server in the other direction (UDP destination port 68).
- For IPv4, clients must be directly-connected to the Firewall Threat Defense and cannot send requests through another relay agent or a router. For IPv6, the Firewall Threat Defense supports packets from another relay server.
- The DHCP clients must be on different interfaces from the DHCP servers to which the Firewall Threat Defense relays requests.
- You cannot enable DHCP Relay on an interface in a traffic zone.

DDNS service configuration

The firewall's DDNS supports only DynDNS service. Hence, ensure that the DDNS is configured with update URL in the following syntax:

https://*username***:***password***@***provider-domain***/***path***?hostname=<h>&myip=<a>**.

Configure the DHCPv4 server

Configure a DHCPv4 server to automatically assign IP addresses, DNS servers, domain names, and other network configuration parameters to DHCP clients on your network.

The DHCPv4 server provides dynamic IP address allocation and network configuration to clients. You can configure global server settings, interface-specific pools, and advanced DHCP options to customize the service for your network requirements.

Before you begin

Follow these steps to configure the DHCPv4 server:

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **DHCP** > **DHCP Server**.

Step 3 Configure these DHCP server options:

- **Ping Timeout**—The amount of time in milliseconds that Firewall Threat Defense device waits to time out a DHCP ping attempt. Valid values range from 10 to 10000 milliseconds. The default value is 50 milliseconds.

  To avoid address conflicts, the Firewall Threat Defense device sends two ICMP ping packets to an address before assigning that address to a DHCP client.

- **Lease Length**—The amount of time in seconds that the client may use its allocated IP address before the lease expires. Valid values range from 300 to 1048575 seconds. The default value is 3600 seconds (1 hour).

- (Routed mode) **Auto-configuration**—Enables DHCP auto configuration on the Firewall Threat Defense device. Auto-configuration enables the DHCP server to provide the DHCP clients with the DNS server, domain name, and WINS server information obtained from a DHCP client running on the specified interface. Otherwise, you can disable auto configuration and add the values yourself in Step 4.

- (Routed mode) **Interface**—Specifies the interface to be used for auto configuration. For a device with virtual routing capability, this interface can only be a global virtual router interface.

Step 4 To override auto-configured settings, do the following:

- Enter the domain name of the interface. For example, your device may be in the Your_Company domain.
- From the drop-down list, choose the DNS servers (primary and secondary) configured for the interface. To add a new DNS server, see Create a network object, on page 2985.
- From the drop-down list, choose the WINS servers (primary and secondary) configured for the interface. To add a new WINS server, see Create a network object, on page 2985.

Step 5 Select **Server**, click **Add**, and configure the following options:

- **Interface**—Choose the interface from the drop-down list. In transparent mode, specify a named bridge group member interface. In routed mode, specify a named routed interface or a named BVI; do not specify the bridge group member interface. Note that each bridge group member interface for the BVI must also be named for the DHCP server to operate.
- **Address Pool**—The range of IP addresses from lowest to highest that is used by the DHCP server. The range of IP addresses must be on the same subnet as the selected interface and cannot include the IP address of the interface itself.
- **Enable DHCP Server**—Enables the DHCP server on the selected interface.

Step 6 Click **OK** to save the DHCP server configuration.

Step 7 (Optional) Select **Advanced**, click **Add**, and specify the type of information you want the option to return to the DHCP client:

- **Option Code**—The Firewall Threat Defense device supports the DHCP options listed in RFC 2132, RFC 2562, and RFC 5510 to send information. All DHCP options (1 through 255) are supported except for 1, 12, 50–54, 58–59, 61, 67, and 82. See DHCPv4 servers, on page 1223for more information on DHCP option codes.

  Note
  The Firewall Threat Defense device does not verify that the option type and value that you provide match the expected type and value for the option code, as defined in RFC 2132. For more information about option codes and their associated types and expected values, see RFC 2132.

- **Type**—DHCP option type. Available options include **IP**, **ASCII**, and **HEX**. If you chose IP, you must add IP addresses in the IP Address fields. If you chose ASCII, you must add the ASCII value in the ASCII field. If you chose HEX, you must add the HEX value in the HEX field.
- **IP Address 1** and **IP Address 2**—The IP address(es) to be returned with this option code. To add a new IP address, see Create a network object, on page 2985.
- **ASCII**—The ASCII value that is returned to the DHCP client. The string cannot include spaces.
- **HEX**—The HEX value that is returned to the DHCP client. The string must have an even number of digits and no spaces. You do not need to use a 0x prefix.

Step 8 Click **OK** to save the option code configuration.

Step 9 Click **Save** on the DHCP page to save your changes.

a) To view DHCP bindings, use **show dhcpd binding**.

Example:

```
> show dhcpd binding
IP Address Client-id         Lease Expiration  Type
10.0.1.100 0100.a0c9.868e.43 84985 seconds     automatic
```

DHCPv6 stateless server configuration

The DHCPv6 stateless server configuration enables the Firewall Threat Defense to provide network information such as the DNS server or domain name to clients using StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature. When clients send Information Request (IR) packets to Firewall Threat Defense, they receive this information.

Create a DHCP IPv6 pool

Create a DHCP IPv6 pool to define the parameters sent in IR messages from the DHCPv6 server to clients.

Create a DHCP IPv6 pool for use with the DHCPv6 server. The DHCPv6 server provides information such as the DNS server or domain name when they clients send Information Request (IR) packets to the Firewall Threat Defense. The DHCP IPv6 pool defines the parameters to sent in IR messages.

This feature is only supported in routed mode. This feature is not supported in clustering or High Availability.

Procedure

Step 1 Choose **Objects** > **Object Management** > **DHCP IPv6 Pool**.

Step 2 Click **Add** ( ).

Step 3 Configure the **DNS Server** and **Domain Name**.

You can either manually define the values and click **Add**, or you can check **Import** to use one or more parameters that the Firewall Threat Defense obtained from the DHCPv6 server on the Prefix Delegation client interface. You can mix and match manually-configured parameters with imported parameters; however, you cannot configure the same parameter manually and also use **Import**.

Figure 319: Manually define values

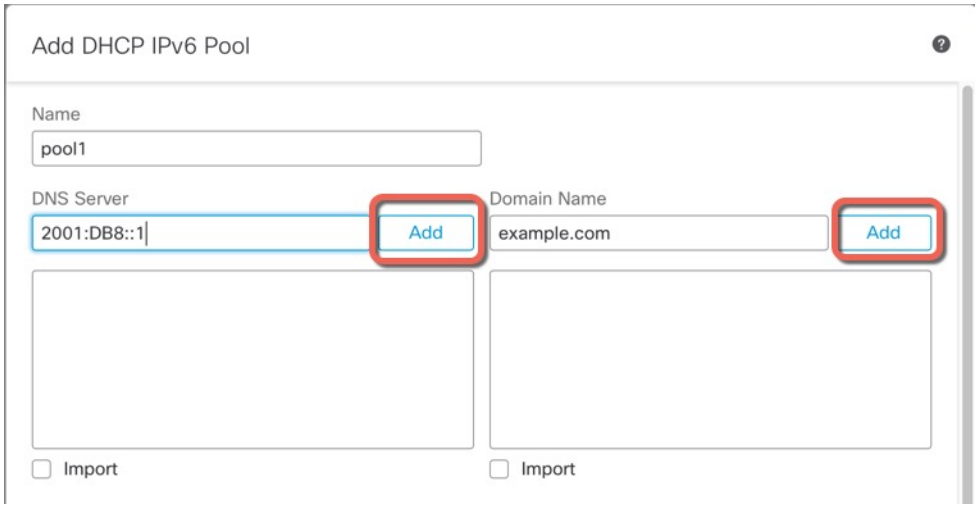
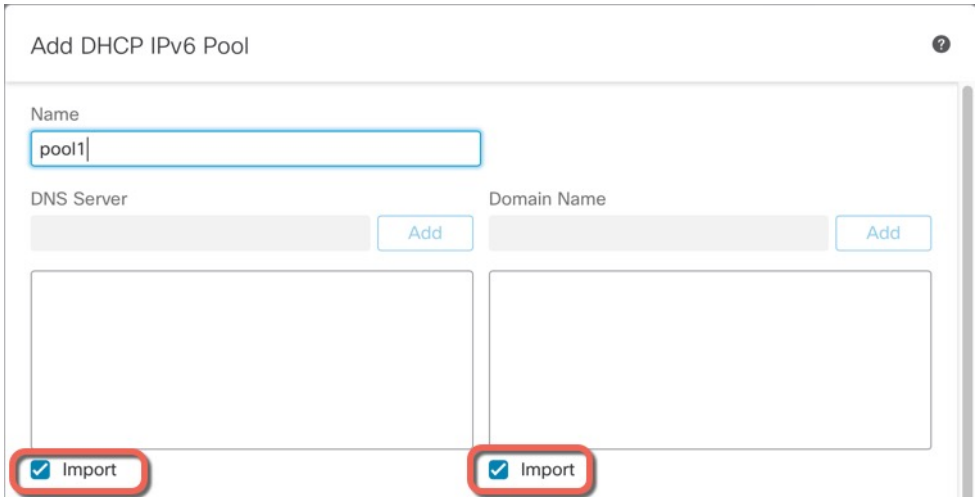


Figure 320: Import values



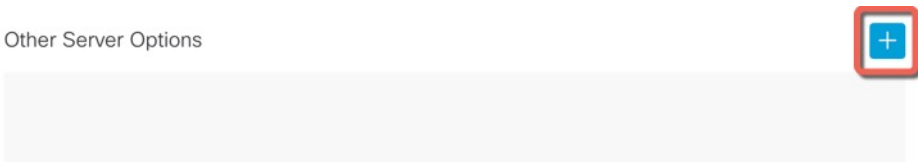
Step 4 Define **Other Server Options**.

You can define the domain name and IP address for these servers:

- NIS
- NISP
- SIP
- SNTP

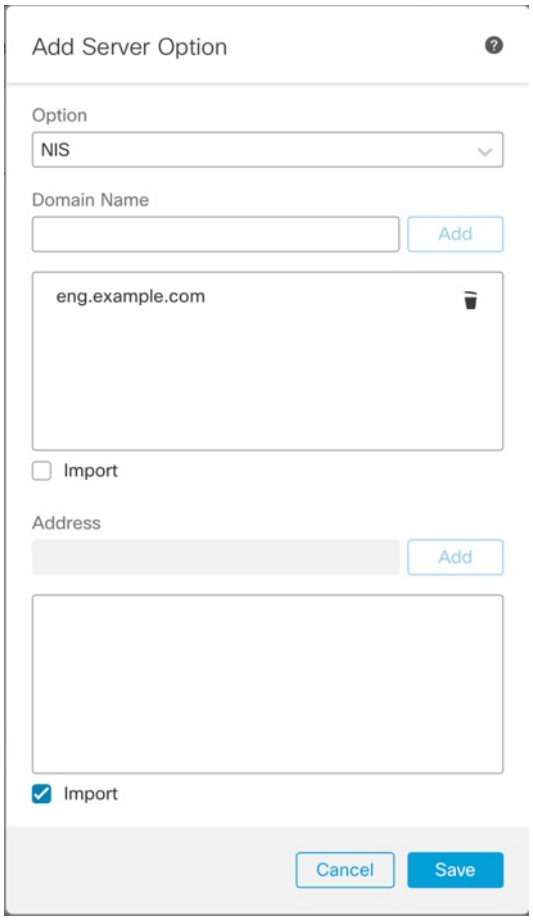
a) Click **Add** ( ).

Figure 321: Other server options



b) Choose the server type under **Option**, and either manually define the **Domain Name** and **Address**, or check **Import**.

Figure 322: Define server domain name and address



Import uses one or more parameters that the Firewall Threat Defense obtained from the DHCPv6 server on the Prefix Delegation client interface. You can mix and match manually-configured parameters with imported parameters; however, you cannot configure the same parameter manually and also use **Import**.

c) Click **Save**.

d) Repeat for each server type.

Step 5 Click **Save**.

Step 6 Use this pool with the DHCPv6 server. See Enable the DHCPv6 stateless server, on page 1232.

Enable the DHCPv6 stateless server

This task enables the DHCPv6 stateless server to provide configuration information to IPv6 autoconfiguration clients without assigning addresses. The server responds to Information Request (IR) packets with DNS server addresses, domain names, and other configuration parameters.

For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature (Enable the IPv6 prefix delegation client, on page 1093), you can configure the Firewall Threat Defense to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the Firewall Threat Defense. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses based on prefixes received in Router Advertisement messages; in other words, based on the prefix that the Firewall Threat Defense received using Prefix Delegation.

This feature is only supported in routed mode. This feature is not supported in clustering or High Availability.

Before you begin

Add a DHCP IPv6 Pool object. See Create a DHCP IPv6 pool, on page 1229. This object defines the server parameters included in the IR messages.

Follow these steps to enable the DHCPv6 stateless server:

Procedure

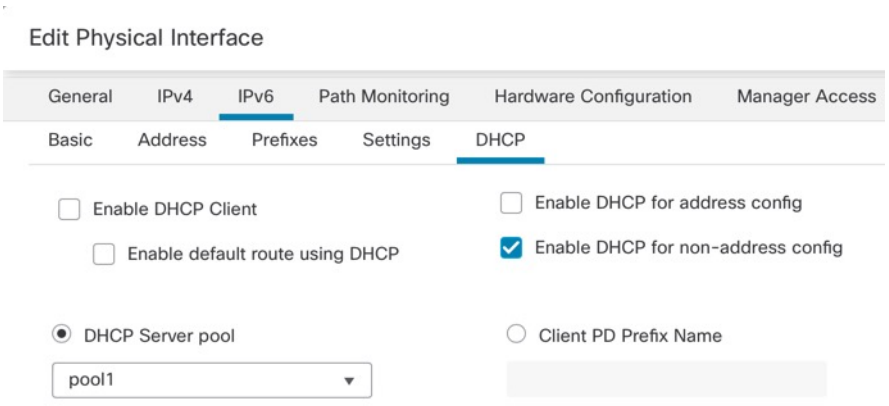
Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **IPv6** page, and then click **DHCP**.

Step 4 Click **DHCP Server Pool**, and choose the object you created earlier.

Figure 323: Enable the DHCPv6 server



Step 5 Check **Enable DHCP for non-address config** to inform SLAAC clients about the DHCPv6 server.

This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure the DHCP relay agent

Configure a DHCP relay agent to forward DHCP requests received on an interface to one or more DHCP servers when DHCP clients and servers are on different network segments.

You can configure a DHCP relay agent to forward DHCP requests received on an interface to one or more DHCP servers. DHCP clients use UDP broadcasts to send their initial DHCPDISCOVER messages because they do not have information about the network to which they are attached. If the client is on a network segment that does not include a server, UDP broadcasts normally are not forwarded by the Firewall Threat Defense device because it does not forward broadcast traffic.

You can remedy this situation by configuring the interface of the Firewall Threat Defense device that is receiving the broadcasts to forward DHCP requests to a DHCP server on another interface.

Note DHCP Relay is not supported in transparent firewall mode.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **DHCP** > **DHCP Relay**.

Step 3 In the **IPv4 Relay Timeout** and **IPv6 Relay Timeout** fields, enter the amount of time in seconds that the Firewall Threat Defense device waits to time out the DHCP relay agent. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.

The timeout is for address negotiation through the local DHCP Relay agent.

Step 4 (Optional) Check **Trust All Information** to set all client interfaces as trusted.

You can configure interfaces as trusted interfaces to preserve DHCP Option 82. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the Firewall Threat Defense DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the Firewall Threat Defense will drop that packet by default. You can preserve Option 82 and forward the packet by identifying an interface as a trusted interface.

Step 5 On **DHCP Relay Agent**, click **Add**, and configure these options:

- **Interface**—The interface connected to the DHCP clients.
- **Enable IPv4 Relay**—Enables IPv4 DHCP Relay for this interface.
- **Set Route**—(For IPv4) Changes the default gateway address in the DHCP message from the server to that of the Firewall Threat Defense device interface that is closest to the DHCP client, which relayed the original DHCP request. This action allows the client to set its default route to point to the Firewall Threat Defense device even if the DHCP server specifies a different router. If there is no default router option in the packet, the Firewall Threat Defense device adds one containing the interface address.
- **Enable IPv6 Relay**—Enables IPv6 DHCP Relay for this interface.

Step 6 Click **OK** to save the DHCP relay agent changes.

Step 7 On **DHCP Servers**, click **Add**, and configure these options:

Add the IPv4 and IPv6 server addresses as separate entries, even if they belong to the same server.

- **Server**—The IP address of the DHCP server. Chose an IP address from the drop-down list. To add a new one, see Create a network object, on page 2985
- **Interface**—The interface to which the specified DHCP server is attached. The DHCP Relay agent and the DHCP server cannot be configured on the same interface.

Step 8 Click **OK** to save the DHCP server changes.

Step 9 Click **Save** on the DHCP page to save your changes.

Configure Dynamic DNS

When an interface uses DHCP IP addressing, the assigned IP address can change when the DHCP lease is renewed. When the interface needs to be reachable using a fully qualified domain name (FQDN), the IP address change can cause the DNS server resource records (RRs) to become stale. Dynamic DNS (DDNS) provides a mechanism to update DNS RRs whenever the IP address or hostname changes. You can also use DDNS for static or PPPoE IP addressing.

DDNS updates these RRs on the DNS server: the A RR includes the name-to-IP address mapping, while the PTR RR maps addresses to names.

The Firewall Threat Defense supports these DDNS update methods:

- Standard DDNS—The standard DDNS update method is defined by RFC 2136.

  With this method, the Firewall Threat Defense and the DHCP server use DNS requests to update the DNS RRs. The Firewall Threat Defense or DHCP server sends a DNS request to its local DNS server for information about the hostname and, based on the response, determines the main DNS server that owns the RRs. The Firewall Threat Defense or DHCP server then sends an update request directly to the main DNS server. See these typical scenarios.

  - The Firewall Threat Defense updates the A RR, and the DHCP server updates the PTR RR.

    Typically, the Firewall Threat Defense "owns" the A RR, while the DHCP server "owns" the PTR RR, so both entities need to request updates separately. When the IP address or hostname changes,

the Firewall Threat Defense sends a DHCP request (including the FQDN option) to the DHCP server to inform it that it needs to request a PTR RR update.

- The DHCP server updates both the A and PTR RR.

  Use this scenario if the Firewall Threat Defense does not have the authority to update the A RR. When the IP address or hostname changes, the Firewall Threat Defense sends a DHCP request (including the FQDN option) to the DHCP server to inform it that it needs to request an A and PTR RR update.

You can configure different ownership depending on your security needs and the requirements of the main DNS server. For example, for a static address, the Firewall Threat Defense should own the updates for both records.

- Web—The Web update method uses the DynDNS Remote API specification (https://help.dyn.com/remote-access-api/).

  With this method when the IP address or hostname changes, the Firewall Threat Defense sends an HTTP request directly to a DNS provider with which you have an account.

The **DDNS** page also supports setting DHCP server settings relating to DDNS.

Note DDNS is not supported on the BVI or bridge group member interfaces.

Before you begin

- Configure a DNS server group on **Objects** > **Object Management** > **DNS Server Group**, and then enable the group for the interface on **Devices** > **Platform Settings**, create or edit a Threat Defense policy and click **DNS**. See Configure DNS server settings, on page 1246.
- Configure the device hostname. You can configure the hostname when you perform the Firewall Threat Defense initial setup, or by using the **configure network hostname** command. If you do not specify the hostname per interface, then the device hostname is used.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **DHCP** > **DDNS**.

Step 3 Standard DDNS method: Configure a DDNS update method to enable DNS requests from the Firewall Threat Defense.

You do not need to configure a DDNS update method if the DHCP server will perform all requests.

a) On **DDNS Update Methods**, click **Add**.
b) Set the **Method Name**.
c) Click **DDNS**.
d) (Optional) Configure the **Update Interval** between DNS requests. By default when all values are set to 0, update requests are sent whenever the IP address or hostname changes. To send requests regularly, set the **Days** (0-364), **Hours**, **Minutes**, and **Seconds**.

e) Set the **Update Records** you want the Firewall Threat Defense to update.

This setting only affects the records you want to update directly from the Firewall Threat Defense; to determine the records you want the DHCP server to update, configure the DHCP client settings per interface or globally. See, Step 5, on page 1236.

- **Not Defined**—Disables DNS updates from the Firewall Threat Defense.
- **Both A and PTR Records**—Sets the Firewall Threat Defense to update both A and PTR RRs. Use this option for static or PPPoE IP addressing.
- **A Records**—Sets the Firewall Threat Defense to update the A RR only. Use this option if you want the DHCP server to update the PTR RR.

f) Click **OK**.

g) Assign this method to the interface in Step 5, on page 1236.

Step 4 Web method: Configure a DDNS update method to enable HTTP update requests from the Firewall Threat Defense.

a) On **DDNS Update Methods**, click **Add**.

b) Set the **Method Name**.

c) Click **Web**.

d) Set the **Web Update Type** to update IPv4, IPv6, or both types of addresses.

e) Set the **Web URL**. Specify the update URL. Check with your DNS provider for the URL required.

Use this syntax:

https://*username***:***password***@***provider-domain***/***path***?hostname=<h>&myip=<a>**

Example:

https://jcrichton:pa$$w0rd17@domains.example.com/nic/update?hostname=<h>&myip=<a>

f) (Optional) Configure the **Update Interval** between DNS requests. By default when all values are set to 0, update requests are sent whenever the IP address or hostname changes. To send requests regularly, set the **Days** (0-364), **Hours**, **Minutes**, and **Seconds**.

g) Click **OK**.

h) Assign this method to the interface in Step 5, on page 1236.

i) The web type method for DDNS also requires you to identify the DDNS server root CA to validate the DDNS server certificate for the HTTPS connection. See, Step 9, on page 1238.

Step 5 Configure interface settings for DDNS, including setting the update method, DHCP client settings, and the hostname for this interface.

a) On **DDNS Interface Settings**, click **Add**.

b) Choose the **Interface** from the drop-down list.

c) Choose the **Method Name** that you created on the **DDNS Update Methods** page.

(Standard DDNS method) You do not need to assign a method if you want the DHCP server to perform all updates.

d) Set the **Host Name** for this interface.

If you do not set the hostname, the device hostname is used. If you do not specify an FQDN, then the default domain from the DNS server group is appended (for static or PPPoE IP addressing) or the domain name from the DHCP server is appended (for DHCP IP addressing).

e) Standard DDNS method: Configure the **DHCP Client requests DHCP server to update requests** to determine which records you want the DHCP server to update.

The Firewall Threat Defense sends DHCP client requests to the DHCP server. Note that the DHCP server must also be configured to support DDNS. The server can be configured to honor the client requests, or it can override the client (in which case, it will reply to the client so the client does not also try to perform updates that the server is performing).

For static or PPPoE IP addressing, these settings are ignored.

Note
You can also set these values globally for all interfaces on the **DDNS** page. The per-interface settings take precedence over the global settings.

- **Not Selected**—Disables DDNS requests to the DHCP server. Even if the client does not request DDNS updates, the DHCP server can be configured to send updates anyway.
- **No Update**—Requests the DHCP server not to perform updates. This setting works in conjunction with a DDNS update method with **Both A and PTR Records** enabled.
- **Only PTR**—Requests that the DHCP server perform the PTR RR update. This setting works in conjunction with a DDNS update method with **A Records** enabled. Use this setting to enable DHCP option 81.
- **Both A and PTR Records**—Requests that the DHCP server perform both A and PTR RR updates. This setting does not require a DDNS update method to be associated with the interface.

f) Click **OK**.

Note
The **Dynamic DNS Update** settings relate to DHCP server settings when you enable a DHCP server on the Firewall Threat Defense. See, Step 6, on page 1237 for more information.

Step 6 If you enable the DHCP server on an Firewall Threat Defense, you can configure DHCP server settings for DDNS.

To enable the DHCP server, see Configure the DHCPv4 server, on page 1227). You can configure the server behavior when DHCP clients use the standard DDNS update method. If the server performs any updates, then if the client lease expires (and is not renewed), the server will request that the DNS server remove the RRs for which it was responsible.

a) You can configure server settings globally or per interface. For global settings, refer to the main **DDNS** page. For per-interface settings, refer to the **DDNS Interface Settings** page. Interface settings take precedence over global settings.
b) Configure which DNS RRs you want the DHCP server to update under **Dynamic DNS Update**.

- **Not Selected**—DDNS updates are disabled, even if the client requests them.
- **Only PTR**—Enables DDNS updates. If you enable the **Override DHCP Client Requests** setting, then the server will only update the PTR RR. Otherwise, the server will update RRs that the client requests. If the client does not send an update request with the FQDN option, the server will request an update for both A and PTR RRs using the hostname discovered in DHCP option 12.
- **Both A and PTR Records**—Enables DDNS updates. If you enable the **Override DHCP Client Requests** setting, then the server will update both the A and PTR RRs. Otherwise, the server will update RRs that the client requests. If the client does not send an update request with the FQDN

option, the server will request an update for both A and PTR RRs using the hostname discovered in DHCP option 12.

c) To override the update actions requested by the DHCP client, check **Override DHCP Client Requests**.

The server will reply to the client that the request was overridden, so the client does not also try to perform updates that the server is performing.

Step 7 (Optional) Configure general DHCP client settings. These settings are not related to DDNS, but are related to how the DHCP client behaves.

a) On the **DDNS** page, check **Enable DHCP Client Broadcast** to request that the DHCP server broadcast the DHCP reply (DHCP option 1).

b) To force a MAC address to be stored inside a DHCP request packet for option 61 instead of the default internally generated string, on **DDNS** > **DHCP Client ID Interface**, choose the interface from the **Available Interfaces** list, and then click **Add** to move it to the **Selected Interfaces** list.

Some ISPs expect option 61 to be the interface MAC address. If the MAC address is not included in the DHCP request packet, then an IP address will not be assigned. This setting does not directly relate to DDNS, but is a general DHCP client setting.

Step 8 Click **Save** on the Device page to save your changes.

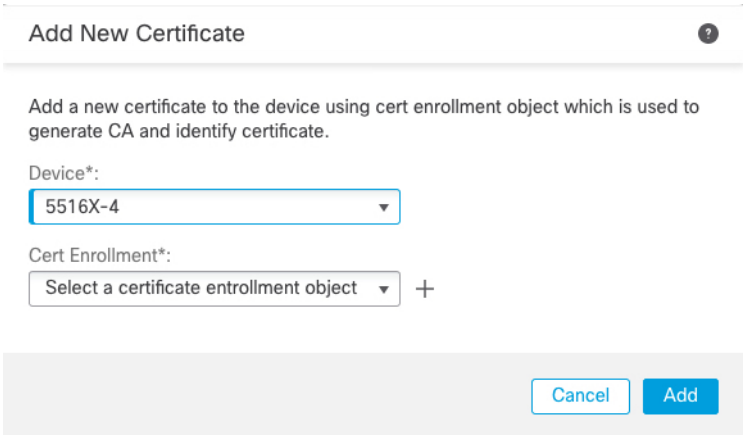
Step 9 The Web method for DDNS also requires you to identify the DDNS server root CA to validate the DDNS server certificate for the HTTPS connection.

This example shows how to add a DDNS server's CA as a trustpoint.

a) Obtain the DDNS server CA certificate. This procedure shows a manual import using PEM format, but you can also use PKCS12.

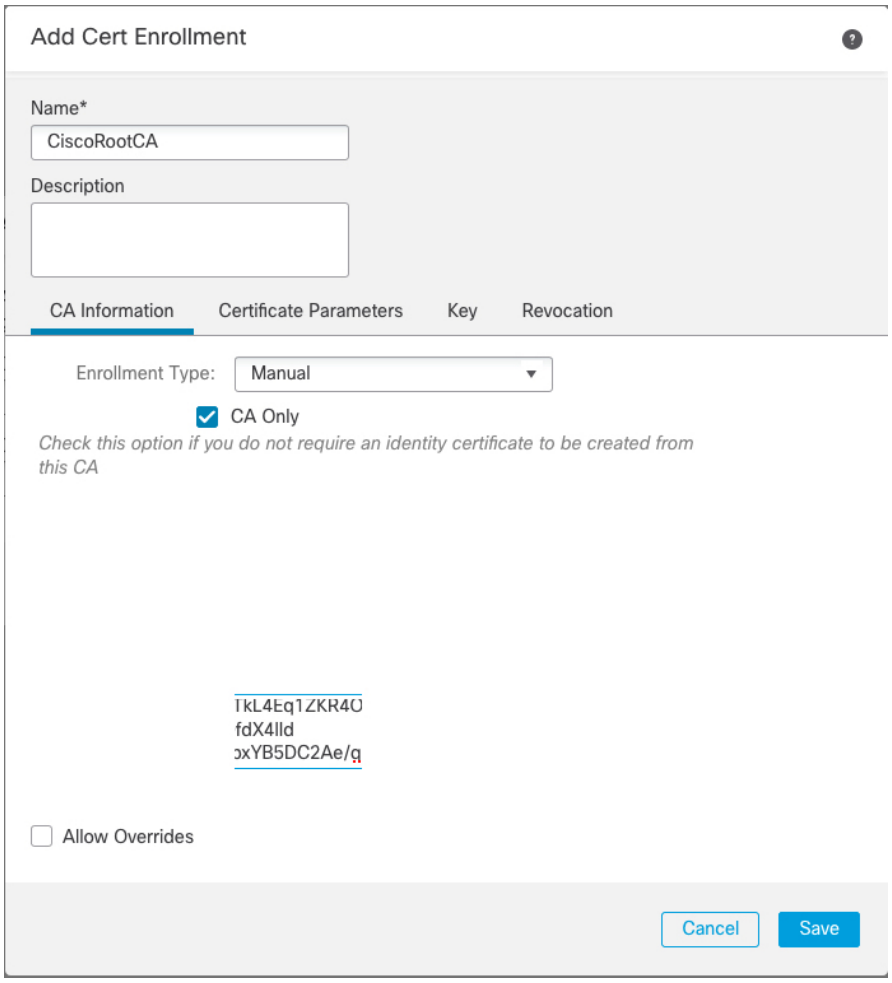
b) In Cloud-Delivered Firewall Management Center, choose **Devices** > **Certificates**, and click **Add**.

c) Select a **Device**, and click **Add** (+).



The **Add Cert Enrollment** dialog box appears.

d) Enter these fields, and click **Save**:



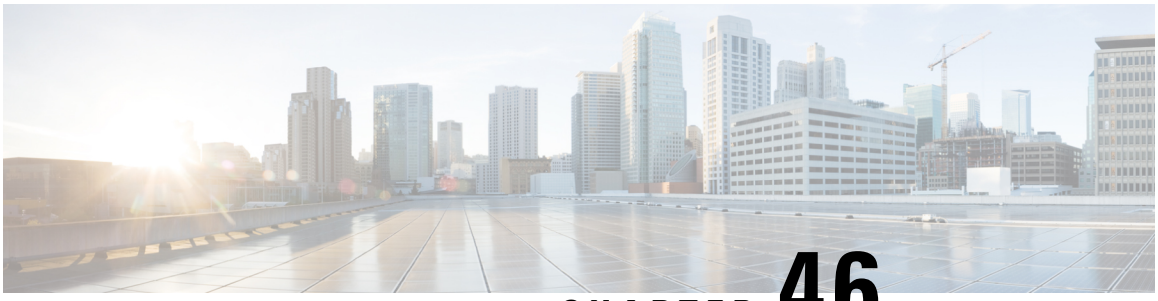
- Enter a **Name**.
- Choose **Enrollment Type** > **Manual**.
- Click **CA Only**.
- Paste in the CA text from step 9.a, on page 1238.

e) Click **Save**.

History for DHCP and DDNS

This reference provides the release history for DHCP and DDNS features, including version requirements and implementation details for Secure Firewall Management Center and Threat Defense devices.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Configure DHCP relay trusted interfaces from the Cloud-Delivered Firewall Management Center web interface.	20240203	Any	You can now use the Cloud-Delivered Firewall Management Center web interface to configure interfaces as trusted interfaces to preserve DHCP Option 82. If you do this, these settings override any existing FlexConfigs, although you should remove them. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the Firewall Threat Defense DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then Firewall Threat Defense will drop that packet by default. You can preserve Option 82 and forward the packet by identifying an interface as a trusted interface. New/modified screens: **Devices** > **Device Management** > **Add/Edit Device** > **DHCP** > **DHCP Relay** See: Configure the DHCP Relay Agent
DHCPv6 Stateless Server	20221213	7.3.0	The Firewall Threat Defense now supports a light DHCPv6 stateless server when using the DHCPv6 Prefix Delegation client. The Firewall Threat Defense provides other information such as the domain name to SLAAC clients when they send Information Request (IR) packets to the Firewall Threat Defense. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. New/modified screens: • **Devices** > **Device Management** > **Interfaces** > **Add/Edit Interfaces** > **IPv6** > **DHCP** • **Objects** > **Object Management** > **DHCP IPv6 Pool** New/modified commands: **show ipv6 DHCP**



CHAPTER 46

Platform Settings

Platform settings for Firewall Threat Defense devices configure a range of unrelated features whose values you might want to share among several devices. Even if you want different settings per device, you must create a shared policy and apply it to the desired device.

Note In 7.4 and later, the Management and Diagnostic interfaces are merged. If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, then you must use separate Platform Settings policies for merged and non-merged devices (7.3 and earlier, and some upgraded 7.4 devices).

- Platform settings policies, on page 1242
- List requirements and prerequisites for platform settings policies, on page 1242
- Manage platform settings policies, on page 1243
- Configure ARP inspection, on page 1244
- Configure a banner, on page 1245
- Configure DNS server settings, on page 1246
- Configure external authentication, on page 1249
- Enable a virtual-router-aware interface for platform external authentication, on page 1255
- Configure fragment settings, on page 1256
- Configure HTTP access, on page 1257
- Configure ICMP access, on page 1258
- Configure NetFlow, on page 1260
- Configure Secure Shell (SSH) access, on page 1262
- Configure an SMTP server, on page 1264
- Configure SNMP for a security device, on page 1264
- Configure SSL settings, on page 1278
- Syslog, on page 1283
- Configure timeouts, on page 1301
- Synchronize time on FTD devices, on page 1303
- Configure a time zone for a device, on page 1304
- Configure UCAPL/CC compliance, on page 1305
- Configure a performance profile, on page 1306

Platform settings policies

A platform settings policy is a shared set of features or parameters that

- define aspects of managed devices that are likely to be similar to other managed devices in a deployment, such as time settings and external authentication,
- enable configuration of multiple managed devices at once, providing consistency and streamlining management efforts, and
- apply changes to all managed devices where the policy is used, requiring creation and application of shared policies even for device-specific settings.

Platform settings policy reference information

Use platform settings policies to enforce security requirements, such as displaying a login banner message, and to manage device configurations efficiently.

- You can set a “No Unauthorized Use” message for appliance logins using platform settings policies.
- You can apply multiple platform settings policies to a single Cloud-Delivered Firewall Management Center to accommodate different mail relay hosts or access lists.

Platform settings policy example

Your organization’s security policies may require that appliances display a “No Unauthorized Use” message when a user logs in. With platform settings, you can set the login banner once in a platform settings policy.

Multiple platform settings policy example

If your organization has different mail relay hosts that you use for different circumstances or if you want to test different access lists, you can create several platform settings policies and switch between them, rather than editing a single policy.

List requirements and prerequisites for platform settings policies

You must meet specific requirements and prerequisites before you configure Platform Settings Policies, including supported domains and user roles.

Supported domains

Any

User roles

Admin

Access Admin

Network Admin

Manage platform settings policies

Manage platform settings policies to control device configurations and assignments within your environment. This task allows you to create, edit, or delete policies for supported device types.

Use the **Platform Settings** page (**Devices** > **Platform Settings**) to manage platform settings policies. This page indicates the type of device for each policy. The **Status** column shows the device targets for the policy.

Procedure

Step 1 Choose **Devices** > **Platform Settings**.

Step 2 For an existing policy, you can **Copy** ( ), **Edit** ( ), or **Delete** ( ) the policy.

Caution
You should not delete a policy that is the last-deployed policy on any of its target devices, even if it is out of date. Before you delete the policy completely, it is a good practice to deploy a different policy to those targets.

Step 3 To create a new policy, click **New Policy**.

a) Choose a device type from the drop-down list:

- **Firepower Settings** to create a shared policy for managed Classic devices.
- **Threat Defense Settings** to create a shared policy for managed Firewall Threat Defense devices.

b) Enter a **Name** for the new policy and optionally, a **Description**.

c) Optionally, choose the **Available Devices** where you want to apply the policy and click **Add** (or drag and drop) to add the selected devices. You can enter a search string in the **Search** field to narrow the list of devices.

d) Click **Save**.

- When you have selected a device and if the device is already applied with a policy, the **Confirm PFS Overwrite** dialog appears. To override the applied policy, click **OK**. The selected device is applied with the new policy.
- The system creates the policy and opens it for editing.

Step 4 To change the target devices for a policy, click **Edit** ( ) next to the platform settings policy that you want to edit.

a) Click **Policy Assignment**.

b) To assign a device, high-availability pair, or device group to the policy, select it in the **Available Devices** list and click **Add**. You can also drag and drop.

c) To remove a device assignment, click **Delete** ( ) next to a device, high-availability pair, or device group in the **Selected Devices** list.

d) Click **OK**.

After completing these steps, platform settings policies are created, edited, or deleted as needed, and device assignments are updated accordingly.

What to do next

- Deploy configuration changes.

Configure ARP inspection

Configure ARP inspection to ensure that only valid ARP packets are accepted on your network, reducing the risk of ARP spoofing and related attacks.

- Apply ARP inspection policies to specific interfaces and security zones.
- Add static ARP entries for enhanced control and security.

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

When you enable ARP inspection, the Firewall Threat Defense device compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the Firewall Threat Defense device drops the packet.
- If the ARP packet does not match any entries in the static ARP table, then you can set the Firewall Threat Defense device to either forward the packet out all interfaces (flood), or to drop the packet.

Note The dedicated Management interface never floods packets even if this parameter is set to flood.

Use this task when you need to protect your network from ARP spoofing by validating ARP packets and controlling ARP traffic on selected interfaces and zones.

Use ARP inspection in environments where network security is a priority and static ARP entries are needed.

Before you begin

Make sure you have device access and permission to edit platform settings and ARP inspection policies.

- Have the list of interfaces and security zones where ARP inspection will be applied.

- Prepare the static ARP entries that you intend to add.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **ARP Inspection**.

Step 3 Add entries to the ARP inspection table.

a) Click **Add** to create a new entry, or click **Edit** if the entry already exists.

b) Select the desired options.

- **Inspect Enabled**—To perform ARP inspection on the selected interfaces and zones.
- **Flood Enabled**—Floods ARP requests that do not match static ARP entries to all interfaces except the originating or dedicated management interface. By default, this behavior is enabled.

  If you choose not to flood ARP requests, only requests that match static ARP entries are allowed.
- **Security Zones**—Add the zones that contain the interfaces on which to perform the selected actions. The zones must be switched zones. For interfaces not in a zone, type the interface name into the field in the Selected Security Zone list and click **Add**. The device applies these rules only if it includes the selected interfaces or zones.

c) Click **OK**.

Step 4 Add static ARP entries according to Add a static ARP entry, on page 1111.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

ARP inspection is enabled on the selected interfaces and security zones. Only valid ARP packets are accepted, and static ARP entries are enforced. The policy is active after deployment.

Configure a banner

Configure a banner to display messages to users when they connect to the device command line interface (CLI). This helps inform users and discourage unauthorized access.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Banner**.

Step 3 Configure the banner.

These are some tips and requirements for banners.

- Use only ASCII characters. You can add line returns (press Enter), but do not use tabs.
- You can dynamically add the hostname or domain name of the device by including the variables **$(hostname)** or **$(domain)**.
- Although banners do not have an absolute length restriction, Telnet or SSH sessions close if the system does not have enough memory to process banner messages.
- Create a banner that discourages unauthorized access by avoiding the words "welcome" or "please," which may seem inviting. The sample banner provided discourages intruders.

```
You have logged in to a secure device.
If you are not authorized to access this device,
log out immediately or risk criminal charges.
```

Step 4 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

The system displays the banner to users when they connect to the device CLI. The changes occur after you deploy the policy.

Configure DNS server settings

Configure DNS server settings to enable hostname resolution for data traffic on managed devices. This task allows you to set up DNS server groups, trusted DNS services, and interface-specific DNS lookups. These configurations support features such as access control rules, NAT, and application detection.

- Set up DNS server groups for resolving FQDNs used in data traffic.
- Configure trusted DNS servers for DNS snooping and application-based PBR.
- Enable DNS lookups on specific interfaces to control routing behavior.

The Domain Name System (DNS) servers are used to resolve hostnames to IP addresses. Data DNS settings apply to services that use FQDNs, such as access control rules and remote access VPN. Special management traffic uses separate DNS settings configured via CLI commands. This procedure applies only to data DNS servers. For management DNS settings, see the CLI **configure network dns servers** and **configure network dns searchdomains** commands.

Managed devices use routing lookups to determine the correct interface for DNS server communications. The system selects the routing table based on the interfaces enabled for DNS. Multiple DNS server groups can be configured to resolve different DNS domains, allowing internal and external DNS resolution as needed.

Trusted DNS servers

You can configure trusted DNS services for DNS snooping. DNS snooping is used to map the application domains to IPs in order to detect the application on the first packet. You can also include the already configured servers in DNS group, DHCP pool, DHCP relay and DHCP client as trusted DNS servers.

For application-based PBR, configure trusted DNS servers. Ensure DNS traffic passes through Firewall Threat Defense in clear text format (because encrypted DNS is not supported) for domain resolution and application detection.

Before you begin

- Ensure you have created one or more DNS server groups. For more information, see Create a DNS server group object, on page 2965.
- Ensure you have created interface objects to connect to the DNS servers.
- Ensure that the managed device has appropriate static or dynamic routes to access the DNS servers.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Threat Defense policy.

Step 2 Click **DNS**, and then click the **DNS Settings** tab.

Step 3 Check **Enable DNS name resolution by device**.

Step 4 Configure the DNS server groups.

a) Do any of these actions in the DNS server group list:

- To add a group, click **Add**. You cannot add another group once there are 30 filter domains configured within the existing list of server groups.
- To edit a group, click **Edit** ( ) next to the group.
- To remove a group, click **Delete** ( ) next to the group. When you remove a group, you do not delete the DNS server group object. You remove it only from this list.

b) When adding or editing a group, configure the following settings, then click **OK**:

- **Select DNS Group**—Select an existing DNS server group object, or click **+** to create a new one.
- **Make as default**—Select this option to make this group the default group. Any DNS resolution request that does not match the filters for other groups will be resolved using the servers in this group.
- **Filter Domains**—For non-default groups only, enter a list of domain names separated by commas, such as example.com,example.net. Do not include spaces.

  The group is used only for DNS resolutions for these domains. You can enter up to 30 separate domains across all groups in this DNS platform settings policy. Each domain name can contain up to 127 characters.

  Note that these filter domains are not related to the default domain name for the group. The filter list can be different from the default domain.

Step 5 (Optional) Enter the **Expiry Entry Timer** and **Poll Timer** values in minutes.

These options apply to FQDNs specified in network objects only and do not apply to FQDNs used in other features.

- **Expire Entry Timer** specifies the minimum time-to-live (TTL) value for the DNS entry, in minutes. If the expiration timer is longer than the entry's TTL, the TTL is increased to the expire entry time value.

If the TTL is longer than the expiration timer, the expire entry time value is ignored. Upon expiration, the entry is removed from the DNS lookup table. When you remove an entry, the device recompiles the table. Frequent removals can increase the processing load on the device. Some DNS entries can have very short TTL (as short as three seconds); use this setting to virtually extend the TTL. The default is 1 minute. The range is 1 to 65535 minutes.

For systems running 7.0 or earlier, the expiration time is actually added to the TTL; it does not specify a minimum value.

- **Poll Timer** specifies the time limit after which the device queries the DNS server to resolve the FQDN defined in a network object. An FQDN is resolved periodically when either the poll timer has expired or the TTL of the resolved IP entry has expired, whichever occurs first.

Step 6 Enable DNS lookups on all interfaces or on specific interfaces. These choices also affect which routing tables are used.

Enabling DNS lookups on an interface is not the same as specifying the source interface for lookups. The Firewall Threat Defense always uses a route lookup to determine the source interface. Management-only interfaces other than the dedicated Management interface cannot be used.

- No interfaces selected—This setting enables DNS lookups on all interfaces. The Firewall Threat Defense checks the data routing table only.
- Specific interfaces selected but not the **Enable DNS Lookup via diagnostic/management interface also** option—Enables DNS lookups on the specified interfaces. The Firewall Threat Defense checks the data routing table only.
- Specific interfaces selected plus the **Enable DNS Lookup via diagnostic/management interface also** option—Enables DNS lookups on the specified interfaces and the Management interface. The Firewall Threat Defense checks the data routing table, and if no route is found, falls back to the management-only routing table.
- Only the **Enable DNS Lookup via diagnostic/management interface also** option—Enables DNS lookups on Management. The Firewall Threat Defense checks only the management-only routing table.

Step 7 Configure Trusted DNS servers.

a) Click the **Trusted DNS Servers** tab.
b) By default, the existing DNS servers that are configured in DHCP pool, DHCP relay, DHCP client, or DNS server group are included as trusted DNS servers. If you want to exclude any of them, uncheck the appropriate check boxes.
c) To add trusted DNS servers, under **Specify DNS Servers**, click **Edit**.
d) In the **Select DNS Servers** dialog box, either choose a host object as the trusted DNS server or directly specify the IP address of the trusted DNS server:

1. To choose existing host objects, under **Available Host Objects**, select the required host object and click **Add** to include it to **Selected DNS Servers**. For information on adding the host objects, see Create a network object, on page 2985.
2. To directly provide the IP address (IPv4 or IPv6) of the trusted DNS server, enter the address in the given text field, and click **Add** to include it to **Selected DNS Servers**.
3. Click **Save**. The added DNS servers are displayed in the **Trusted DNS Servers** page.

Note
You can configure a maximum of 12 DNS servers per policy.

To search for a DNS server that was added, using either the host name or the IP address, use the search field under **Specify DNS Servers**.

Step 8 Click **Save**.

What to do next

To use FQDN objects for access control rules, create an FQDN network object which can then be assigned to an access control rule. For instructions, see Create a network object, on page 2985.

Configure external authentication

Configure external authentication to allow management users to be authenticated using an LDAP or RADIUS server, as specified in an external authentication object. Centralized user management also enhances security for device access.

Note You must have administrator privileges to perform this task.

When you enable external authentication for management users, the Firewall Threat Defense verifies the user credentials with an LDAP or RADIUS server as specified in an external authentication object.

Sharing External Authentication Objects

External authentication objects can be used by the Cloud-Delivered Firewall Management Center and Firewall Threat Defense devices. You can share the same object between the Cloud-Delivered Firewall Management Center and devices. Alternatively, create separate objects. Note that the Firewall Threat Defense supports defining users on the RADIUS server, while the Cloud-Delivered Firewall Management Center requires you to predefine the user list in the external authentication object. You can choose to use the predefined list method for the Firewall Threat Defense, but if you want to define users on the RADIUS server, you must create separate objects for the Firewall Threat Defense and the Cloud-Delivered Firewall Management Center.

Note The timeout range is different for the Firewall Threat Defense and the Cloud-Delivered Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1 to 30 seconds for LDAP, and 1 to 300 seconds for RADIUS). If you set the timeout to a higher value, the Firewall Threat Defense external authentication configuration will not work.

Assigning External Authentication Objects to Devices

For the Cloud-Delivered Firewall Management Center, enable the external authentication objects directly on **System** (⚙) > **Users** > **External Authentication**; this setting only affects Cloud-Delivered Firewall Management Center usage, and it does not need to be enabled for managed device usage. For Firewall Threat Defense devices, you must enable the external authentication object in the platform settings that you deploy to the devices, and you can only activate one external authentication object per policy. An LDAP object with CAC authentication enabled cannot also be used for CLI access. Be sure that both the Firewall Threat Defense and the Cloud-Delivered Firewall Management Center can reach the LDAP server, even if you are not sharing

the object. The Cloud-Delivered Firewall Management Center is essential to retrieving the user list and downloading it to the device.

Firewall Device Supported Fields

Only a subset of fields in the external authentication object are used for Firewall Threat Defense SSH access. If you fill in additional fields, they are ignored. If you also use this object for the Cloud-Delivered Firewall Management Center, those fields will be used. This procedure only covers the supported fields for the Firewall Threat Defense. For other fields, see *Configure External Authentication for the Cloud-Delivered Firewall Management Center* in the *Cisco Secure Firewall Management Center Administration Guide*.

Usernames

Usernames must be Linux-valid usernames and be lowercase only, using alphanumeric characters plus period (.) or hyphen (-). Special characters such as at sign (@) and slash (/) are not supported. You cannot add these default internal users for external authentication—**root, daemon, bin, sys, sync, games, man, lp, mail, news, uucp, proxy, www-data, backup, list, irc, gnats, ntp, lldpd, bind, dhcp, qemu, rpcuser, rpc, sshd, radiusd, messagebus, admin, mysql, www, sfrna, snorty, sfsnort, sfremediation, nobody, samdme, extuser, p2puser, rloginUser**. You can only add external users (as part of the External Authentication object) in the Cloud-Delivered Firewall Management Center; you cannot add them at the CLI. Note that internal users can only be added at the CLI, not in the Cloud-Delivered Firewall Management Center.

If you previously configured the same username for an internal user using the **configure user add** command, the Firewall Threat Defense first checks the password against the internal user, and if that fails, it checks the AAA server. Note that you cannot later add an internal user with the same name as an external user; only pre-existing internal users are supported. For users defined on the RADIUS server, be sure to set the privilege level to be the same as any internal users; otherwise you cannot log in using the external user password.

Privilege Level

LDAP users always receive Config privileges. You can define RADIUS users as either config or basic users.

Before you begin

- SSH access is enabled by default on the management interface. To enable SSH access on data interfaces, see Configure Secure Shell (SSH) access, on page 1262.
- Inform RADIUS users of the following behavior to set their expectations appropriately:
  - The first time an external user logs in, Firewall Threat Defense creates the required structures but cannot simultaneously create the user session. The user simply needs to authenticate again to start the session. The user will see a message similar to this: "New external username identified. Please log in again to start a session."
  - If the user's Service-Type attribute is not defined or incorrectly configured in the RADIUS server and you use the RADIUS-defined users for authentication, the user will see a message similar to this: "Your username is not defined with a service type that is valid for this system. You are not authorized to access the system?".

    In some cases, the SSH clients close the CLI window on an unsuccessful SSH connection, even before displaying the failure message. Hence, ensure that the user's Service-Type attribute is correctly defined in the RADIUS server.
  - Similarly, if the user’s Service-Type authorization was changed since the last login, the user will need to re-authenticate. The user will see a message similar to this: "Your authorization privilege has changed. Please log in again to start a session."

- The server used for external authentication must be reachable by the Cloud-Delivered Firewall Management Center; otherwise, users' authentication will fail.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **External Authentication**, and then click the **Manage External Authentication Server** link.

You can also open the External Authentication screen by clicking **System** > **Users** > **External Authentication**.

Step 3 Click **Add External Authentication Object**.

Step 4 Configure an LDAP Authentication Object.

Step 5 Configure a RADIUS Authentication Object.

Step 6 Return to **System** (⚙) > **Users** > **External Authentication**.

Step 7 Click **Refresh** (↻) to view any newly-added objects.

For LDAP when you specify SSL or TLS encryption, you must upload a certificate for the connection; otherwise, the server will not be listed on this window.

Step 8 Click **Slider enabled** (⬭) next to the External Authentication object you want to use. You can only enable one object.

Step 9 Click **Save**; deploy configuration changes (see Deploy configuration changes, on page 202).

After completing this task, external authentication is enabled for management users. The system verifies user credentials with the configured LDAP or RADIUS server and grants access based on the authentication object and user privileges.

Configure an LDAP External Authentication Object

Configure external authentication to allow management users to be authenticated using an LDAP server.

Note You must have administrator privileges to perform this task.

Procedure

Step 1 In the **External Authentication Object** page, select **LDAP** from the **Authentication Method** drop-down list and enter appropriate values in these fields:

a) Enter a **Name**.

b) (Optional) Enter **Description**.

Step 2 Choose a **Server Type** from the drop-down list.

Step 3 For the **Primary Server** enter these values:

a) Enter a **Host Name/IP Address**.

Note
If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

b) Change the **Port** from the default.

Step 4 Enter the **Backup Server** parameters.

Step 5 Enter **LDAP-Specific Parameters**.

- **Base DN**—Enter the base distinguished name for the LDAP directory you want to access. For example, to authenticate names in the Security organization at the Example company, enter `ou=security,dc=example,dc=com`. Alternatively click **Fetch DNs**, and choose the appropriate base distinguished name from the drop-down list.
- (Optional) **Base Filter**—For example, if the user objects in a directory tree have a `physicalDeliveryOfficeName` attribute and users in the New York office have an attribute value of `NewYork` for that attribute, to retrieve only users in the New York office, enter `(physicalDeliveryOfficeName=NewYork)`.
- **User Name**—Enter a distinguished name for a user who has sufficient credentials to browse the LDAP server. For example, if you are connecting to an OpenLDAP server where user objects have a `uid` attribute, and the object for the administrator in the Security division at our example company has a `uid` value of `NetworkAdmin`, you might enter `uid=NetworkAdmin,ou=security,dc=example,dc=com`.
- **Password** and **Confirm Password**—Enter and confirm the password for the user.
- (Optional) **Show Advanced Options**—Configure the following advanced options.
  - **Encryption**—Click **None**, **TLS**, or **SSL**.

    Note
    If you change the encryption method after specifying a port, you reset the port to the default value for that method. For **None** or **TLS**, the port resets to the default value of 389. If you choose SSL encryption, the port resets to 636.

  - **SSL Certificate Upload Path**—For SSL or TLS encryption, you must choose a certificate by clicking **Choose File**.
  - (Not Used) **User Name Template**—Not used by the Firewall Threat Defense.
  - **Timeout**—Enter the number of seconds before rolling over to the backup connection between 1 and 30. The default is 30.

    Note
    The timeout range is different for the Firewall Threat Defense and the Cloud-Delivered Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1-30 seconds). If you set the timeout to a higher value, the Firewall Threat Defense external authentication configuration will not work.

Step 6 (Optional) Set the **CLI Access Attribute** if you want to use a shell access attribute other than the user distinguished type. For example, on a Microsoft Active Directory Server, use the `sAMAccountName` shell access attribute to retrieve shell access users by typing `sAMAccountName` in the **CLI Access Attribute** field.

Step 7 Set the **CLI Access Filter**.

Choose one of these methods:

- To use the same filter you specified when configuring authentication settings, choose **Same as Base Filter**.
- To retrieve administrative user entries based on attribute value, enter the attribute name, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses. For example, if all network administrators have a `manager` attribute which has an attribute value of `shell`, you can set a base filter of `(manager=shell)`.

The names on the LDAP server must be Linux-valid usernames:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Step 8 Click **Save**.

Step 9 For LDAP, if you later add or delete users on the LDAP server, you must refresh the user list and redeploy the Platform Settings.

a) Choose **System** (⚙) > **Users** > **External Authentication**.

b) Click **Refresh** (↻) next to the LDAP server.

If the user list changed, you will see a message advising you to deploy configuration changes for your device. The Firepower Theat Defense Platform Settings will also show that it is "Out-of-Date on *x* targeted devices."

c) Deploy configuration changes; see Deploy configuration changes, on page 202.

Configure a RADIUS Authentication Object

Configure external authentication to allow management users to be authenticated using a RADIUS server

Note You must have administrator privileges to perform this task.

Procedure

Step 1 In the **External Authentication Object** page, select **RADIUS** from the **Authentication Method** drop-down list and enter appropriate values in these fields:

a) Check the **Enabled Message Authenticator** check box to require the Message-Authenticator attribute in all RADIUS responses, ensuring that every response from the RADIUS server is securely verified by the Firewall Threat Defense.

This feature is enabled by default for new RADIUS servers. We recommend you enable it for existing servers after the upgrade. Disabling message authenticators may expose your firewalls to potential attacks. Ensure that your RADIUS server has the Message-Authenticator configuration.

b) Define users on the RADIUS server using the Service-Type attribute.

This is the list of supported values for the Service-Type attribute:

- Administrator (6)—Provides Config access authorization to the CLI. These users can use all commands in the CLI.
- NAS Prompt (7) or any level other than 6—Provides Basic access authorization to the CLI. These users can use read-only commands, such as **show** commands, for monitoring and troubleshooting purposes.

The names must be Linux-valid usernames:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Alternatively, you can predefine users in the external authentication object (see Step Configure external authentication, on page 1249). To use the same RADIUS server for the Firewall Threat Defense and Cloud-Delivered Firewall Management Center while using the Service-Type attribute method for the Firewall Threat Defense, create two external authentication objects that identify the same RADIUS server: one object includes the predefined **CLI Access Filter** users (for use with the Cloud-Delivered Firewall Management Center), and the other object leaves the **CLI Access Filter** empty (for use with Firewall Threat Defenses).

c) Enter a **Name** and optional **Description**.

Step 2 For the **Primary Server**, enter a **Host Name/IP Address**.

Note
If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

Step 3 (Optional) Change the **Port** from the default.

Step 4 Enter a **RADIUS Secret Key**.

Step 5 (Optional) Enter the **Backup Server** parameters.

Step 6 Enter **RADIUS-Specific Parameters**.

- **Timeout (Seconds)**—Enter the number of seconds before rolling over to the backup connection. The default is 30.
- **Retries**—Enter the number of times the primary server connection should be tried before rolling over to the backup connection. The default is 3.

Step 7 (Optional) Enter a comma-separated list of usernames in the **Administrator CLI Access User List** field. For example, enter `jchrichton, aerynsun, rygel`.

You may want to use the **CLI Access Filter** method for Firewall Threat Defense so you can use the same external authentication object with Firewall Threat Defense and other platform types. Note that if you want to use RADIUS-defined users, you must leave the **CLI Access Filter** empty.

Make sure that these usernames match usernames on the RADIUS server. The names must be Linux-valid usernames:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Note
If you want to only define users on the RADIUS server, you must leave this section empty.

Step 8 Click **Save**.

Enable a virtual-router-aware interface for platform external authentication

Enable a virtual-router-aware data interface for external authentication of the platform to allow Authentication, Authorization, and Accounting (AAA) management traffic to be routed through the data interface instead of the default management interface.

Authentication, Authorization, and Accounting (AAA) for the Firewall Threat Defense device is managed through the management interface of the device. You can also enable virtual-router-aware data interface, data sub-interface, port-channel, or sub port-channel to manage AAA for the Firewall Threat Defense device. When enabled, the AAA route lookup is in the Virtual Routing and Forwarding (VRF) routing domain, and the AAA management traffic is forwarded to the data interfaces. These server configuration are supported when using virtual-router-aware data interfaces for AAA:

- RADIUS or LDAP servers for external authentication
- FQDN, IPv4, or IPv6 server addresses

To use a virtual-router-aware interface for external authentication of a Firewall Threat Defense device, modify its external authentication policy by associating the authentication servers with the virtual-router-aware interface of the device.

Before you begin

- Ensure that you have configured the required Virtual Routing and Forwarding (VRF) interface with a static route for the device. For information about configuring a VRF interface, see Configure a virtual router, on page 1387, and for information about adding a static route, see Add a static route, on page 1361.
- Ensure that security zones or interface groups with a single virtual-router-aware interface exists. For information about creating security zones and interface groups, see Create security zone and interface group objects, on page 1002.
- If the primary authentication server is configured with the FQDN of the server, ensure that the backup authentication server, if configured, is also with its FQDN. In addition, configure the DNS server in the

threat defense device's management interface. For information about the DNS server configuration, see Configure DNS.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **External Authentication** and edit the external authentication policy.

Step 3 In the **External Authentication** dialog box, the available security zone and interface groups are listed. To associate a virtual-router-aware interface with the external authentication servers, select the security zone or interface group having a single virtual-router-aware interface, and then do the following:

a) To associate the interface object with the primary authentication server, click **Add to Primary Server**.

b) (Optional) To associate the interface object with the backup authentication server, click **Add to Backup Server**. If the **Add to Backup Server** button is inactive, it means that a backup server for external authentication is not configured in the device.

Step 4 Click **OK**.

Step 5 Save and deploy the changes.

Configure fragment settings

Configure fragment settings to manage the handling of fragmented IP packets on your Firewall Threat Defense device. These settings help prevent Denial of Service (DoS) attacks and ensure applications that require packet fragmentation operate properly.

- Set limits for the number of fragments per packet and for fragments awaiting reassembly.
- Set the timeout for fragment reassembly.
- Allow fragments only if your network applications require them.

By default, the Firewall Threat Defense device allows up to 24 fragments per IP packet. It also allows up to 200 fragments awaiting reassembly. Applications such as NFS over UDP may require fragmented packets, but attackers often use fragmented packets in Denial of Service (DoS) attacks.

These settings establish the defaults for devices assigned this policy. You can override these settings for specific interfaces on a device by selecting **Override Default Fragment Setting** in the interface configuration. When you edit an interface, you can find the option on **Advanced** > **Security Configuration**. Select **Devices** > **Device Management**, edit a Firewall Threat Defense device, and select **Interfaces** to edit interface properties.

- Allow fragments only if required by your network applications.
- Override the default settings for specific interfaces when your configuration requires it.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Fragment Settings**.

Step 3 Configure the available options. To use the default settings, click **Reset to Defaults**.

- **Size (Block)**—The maximum number of packet fragments from all connections collectively that can be waiting for reassembly. The default value is 200 fragments. (One fragment equals one packet segment.)
- **Chain (Fragment)**—The maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets. Set this option to one to disallow fragments.
- **Timeout (Sec)**—The maximum number of seconds to wait for an entire fragmented packet to arrive. The default is five seconds. If all fragments are not received within this time, all fragments are discarded.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

Deploy the policy to activate the new fragment settings on assigned devices. Monitor device performance and change fragment settings as needed.

Configure HTTP access

You can enable the HTTPS server to provide a health check mechanism for a cloud load balancer, for example, for the Firewall Threat Defense Virtual on AWS using an Application Load Balancer.

Use HTTPS with the Firewall Threat Defense only as described for health checks and management. You cannot configure the Firewall Threat Defense through a web interface in this management mode.

Configure HTTP access only on data interfaces, including interfaces you set as management-only. Use only data interfaces for HTTP access, not the dedicated Management interface. The Management interface is separate from the other interfaces on the device. Use the Management interface only to set up and register your device to the Cloud-Delivered Firewall Management Center. It has a separate IP address and static routing.

Configure HTTPS access as described in this section. An access rule for the host IP address is not required.

Use HTTPS only on a reachable interface. If your HTTPS host is on the outside interface, initiate the management connection directly to it.

Before you begin

- Configure HTTPS and AnyConnect VPN module of Cisco Secure Client on separate TCP ports if both features are required on the same interface. For example, assign a different port to HTTPS, such as port 4443.
- Create network objects to define the hosts or networks that are allowed for HTTPS connections. You can add these objects as part of this procedure. If you plan to use object groups for IP addresses, ensure the groups exist in the rules. Select **Objects** > **Object Management** to configure objects.

Note Use the **any-ipv4** or **any-ipv6** network object group instead of the system-provided **any** object.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **HTTP Access**.

Step 3 Check the **Enable HTTP Server** check box to enable the HTTP server.

Step 4 (Optional) Change the HTTP port. The default is 443.

Step 5 Identify the interfaces and IP addresses that allow HTTP connections.

Allow HTTP connections only from specific interfaces and client IP addresses by configuring this table. You may use network addresses instead of individual IP addresses.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.

b) Configure the rule properties:

- **IP Address**—The network object or group that identifies the hosts or networks you are allowing to make HTTP connections. Choose an object from the drop-down menu, or click + to add a new network object.
- **Available Zones/Interfaces**—Add the zones that contain the interfaces to which you will allow HTTP connections. For interfaces not in a zone, type the interface name into the field under the selected **Selected Zones/Interfaces** list. Then click **Add**. The device applies these rules only if it includes the selected interfaces or zones.

c) Click **OK**.

Step 6 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

After you complete this task, the device enables HTTP access on the specified interfaces and IP addresses. Use secure connections for health checks or management.

Configure ICMP access

By default, you can send ICMP packets to any interface using either IPv4 or IPv6, with these exceptions:

- The Firewall Threat Defense does not respond to ICMP echo requests directed to a broadcast address.
- The Firewall Threat Defense only responds to ICMP traffic sent to the interface that traffic comes in on; you cannot send ICMP traffic through an interface to a far interface.

To protect the device from attacks, you can use ICMP rules to limit ICMP access to interfaces to particular hosts, networks, or ICMP types. ICMP rules function like access rules, where the rules are ordered, and the first rule that matches a packet defines the action.

If you configure any ICMP rule for an interface, an implicit deny ICMP rule is added to the end of the ICMP rule list, changing the default behavior. Thus, if you want to simply deny a few message types, you must include a permit any rule at the end of the ICMP rule list to allow the remaining message types.

We recommend that you always grant permission for the ICMP unreachable message type (type 3). Denying ICMP unreachable messages disables ICMP path MTU discovery, which can halt IPsec and PPTP traffic. Additionally ICMP packets in IPv6 are used in the IPv6 neighbor discovery process.

Before you begin

Make sure the required objects exist before you configure rules. Select **Objects** > **Object Management** to configure objects. Create network objects or groups to specify the hosts or networks, and port objects to specify the ICMP message types you want to control.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **ICMP Access**.

Step 3 Configure ICMP rules.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.

b) Configure the rule properties:

- **Action**—Whether to permit (allow) or deny (drop) matching traffic.
- **ICMP Service**—The port object that identifies the ICMP message type.
- **Network**—The network object or groupthat identifies the hosts or networks whose access you are controlling.
- **Available Zones/Interfaces**—Add the zones that contain the interfaces that you are protecting. If an interface is not in a zone, type the interface name into the field beneath the **Selected Zones/Interfaces** list, and click **Add**. A device applies these rules only if it includes the selected interfaces or zones.

c) Click **OK**.

Step 4 (Optional.) Set rate limits on ICMPv4 Unreachable messages.

- **Rate Limit**—Sets the rate limit of unreachable messages, between 1 and 100 messages per second. The default is 1 message per second.
- **Burst Size**—Sets the burst rate, between 1 and 10. The system sends this number of replies, but subsequent replies are not sent until the rate limit is reached.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

ICMP access rules take effect after you deploy the policy to the assigned devices. Your configuration permits or denies ICMP traffic for the hosts or networks you specified.

Configure NetFlow

Configure NetFlow to enable the collection and export of IP network traffic information from your device interfaces. This allows you to analyze traffic patterns, bandwidth usage, and identify sources and causes of network congestion.

The NetFlow feature enables you to collect IP network traffic information as it enters or exits an interface. The collected traffic information is sent as records to a NetFlow Collector server or NetFlow Analyzer. You can analyze the data from NetFlow and determine information such as source and destination of traffic, class of service, traffic pattern, bandwidth usage, type of traffic, traffic volume, and the causes of congestion.

With the native NetFlow configuration support, the traffic information collection that was enabled through syslogs flow exports has to be disabled.

NetFlow provides the option to configure the flow exporter and collectors along with flow event types that must be monitored.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **NetFlow**.

Step 3 Enable the **Enable Flow Export** toggle to enable NetFlow data export.

Step 4 Configure the general NetFlow parameters that control the frequency of events pushed to the collector.

a) **Active Refresh Interval**—For active connections, specify the time interval (in minutes) between flow-update events.

b) **Delay Flow Create**—Specify the delay (in seconds) before sending a flow-create event. If you do not enter any value, there is no delay and the flow-create event is exported as soon as the flow is created.

c) **Template Timeout Rate**—Specify the time interval (in minutes) at which the template records are sent to the collectors.

Step 5 Click **Add Collector** to configure the collector. See Add a collector in NetFlow, on page 1260.

Step 6 Click **Add Traffic Class** to configure the traffic class. See Add a traffic class to NetFlow, on page 1261.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

After completing this task, NetFlow will be enabled and configured on your device. Traffic data will be exported to the configured collector, allowing you to analyze network traffic patterns and usage.

Add a collector in NetFlow

Use this task when you need to export NetFlow data from Cisco Threat Defense devices to a collector for traffic analysis and monitoring.

This configuration is applicable to Cisco Threat Defense policies with NetFlow enabled.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Choose **NetFlow**.

Step 3 Enable the **Enable Flow Export** toggle to enable NetFlow data export.

Step 4 Click **Add Collector** to configure the collector. You can configure a maximum of five collectors.

Step 5 From the **Host** dropdown list, choose the collector host IP address (IPv4 only) of the NetFlow event collector or server to which the NetFlow packets must be sent. Alternatively, you can click the (+) icon to create a new network host.

Step 6 In the **Port** field, enter the UDP port on the collector to which the NetFlow packets must be sent.

Step 7 From **Available Interfaces or Interface Groups**, choose the interface or interface group through which the collector must be reached. You can choose multiple interfaces or interface groups. The Interface Group object can contain only one interface for a given device. A collector can be reached only through one interface. The object can include a virtual-router-aware interface. Loopback interface is not supported in NetFlow.

You can click the (+) icon to create a new interface group.

Step 8 Click **Add** to add the interfaces you chose. Alternatively, you can also enter the interface name and click **Add** to add the interface.

Step 9 Click **OK**.

Add a traffic class to NetFlow

Use this task when you need to monitor, filter, and export specific types of network traffic using NetFlow on a Threat Defense device. Adding a traffic class helps you focus on relevant traffic for security or operational analysis.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **NetFlow**.

Step 3 Click the **Enable Flow Export** toggle to export NetFlow data.

Step 4 Click **Add Traffic Class** to configure the traffic class.

Step 5 In the **Name** field, enter the name of the traffic class that must match the NetFlow events.

Step 6 In the **Type** field, choose the traffic class to filter the type of traffic you want to capture:

- **Default**—The traffic class that is matched if none of the traffic classes matches the traffic.
- **Access List**—The specific traffic class that must match the traffic captured for the NetFlow events.

Step 7 If you select **Access List** as the **Type**, then select the access list object from the **Access List Object** drop-down list.

Note

You can also click the (+) icon to create a new extended access list object. See Configure an extended ACL object, on page 2949.

Step 8 In **Event Types**, check the checkboxes for the different NetFlow events that you want to capture and send to the collectors.

Step 9 Click **OK**.

Configure Secure Shell (SSH) access

If you enabled Cloud-Delivered Firewall Management Center access on a data interface, such as outside, you should enable SSH on that interface using this procedure. This section describes how to enable SSH connections to one or more *data* interfaces on the Firewall Threat Defense.

Note SSH is enabled by default on the Management interface; however, this screen does not affect Management SSH access.

To use SSH, you do not also need an access rule allowing the host IP address. You only need to configure SSH access according to this section.

Data vs. Management SSH

The Management interface is separate from the other interfaces on the device. It is used to set up and register the device to the Cloud-Delivered Firewall Management Center. SSH for data interfaces shares the internal and external user lists with SSH for the Management interface. Other settings are configured separately. For data interfaces, enable SSH and access lists using this screen. SSH traffic for data interfaces uses the regular routing configuration. Routing for the Management interface configured during setup or at the CLI is not used.

For the Management interface, to configure an SSH access list, see the **configure SSH-access-list** command in the *Cisco Secure Firewall Threat Defense Command Reference*. To configure a static route, see the **configure network static-routes** command. By default, you configure the default route through the Management interface at initial setup.

Interface Access

You can SSH only to a reachable interface(including an interface in a user-defined virtual router); for example, if your SSH host is located on the outside interface, you can only initiate a management connection directly to the outside interface. When you enable SSH in a user-defined virtual router, and you want VPN users to access SSH, be sure to terminate the VPN on the same virtual router. If the VPN is terminated on another virtual router, then you must configure route leaks between the virtual routers.

Ciphers and Key Exchange

SSH supports the following ciphers and key exchange:

- Encryption—aes128-cbc, aes192-cbc, aes256-cbc, aes128-ctr, aes192-ctr, aes256-ctr

- Integrity—hmac-sha2-256
- Key exchange—dh-group14-sha256

SSH Stack

The Firewall Threat Defense uses the CiscoSSH stack, which is based on OpenSSH. CiscoSSH supports FIPS compliance and regular updates, including updates from Cisco and the open source community.

Maximum Attempts

After you make three consecutive failed attempts to log into the CLI using SSH, the device terminates the SSH connection.

Before you begin

- You can configure SSH internal users at the CLI using the **configure user add** command; see Add an internal user at the CLI, on page 183. By default, there is an **admin** user for which you configured the password during initial setup. You can also configure external users on LDAP or RADIUS by configuring **External Authentication** in platform settings. See Configure external authentication, on page 1249.
- You need network objects that define the hosts or networks permitted to make SSH connections to the device. You can add objects during the procedure. If you want to use object groups to identify a group of IP addresses, make sure the required groups already exist in the rules. Select **Objects** > **Object Management** to configure objects.



Note You cannot use the system-provided **any** network object. Instead, use **any-ipv4** or **any-ipv6**.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **SSH Access**.

Step 3 Identify the interfaces and IP addresses that allow SSH connections.

Use this table to limit which interfaces will accept SSH connections, and the IP addresses of the clients who are allowed to make those connections. You can use network addresses rather than individual IP addresses.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.
b) Configure the rule properties:

- **IP Address**—The network object or group that identifies the hosts or networks you are allowing to make SSH connections. Choose an object from the drop-down menu, or click + to add a new network object.
- **Available Zones/Interfaces**—Add the zones that contain the interfaces to which you will allow SSH connections. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zones/Interfaces** list and click **Add**. You can also add loopback interfaces and virtual-router-aware interfaces. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure an SMTP server

Configure an SMTP server to enable email alerts in Syslog. Ensure the source email address configured for Syslog is a valid account on the SMTP servers.

Before you begin

Ensure that the network objects for the primary and secondary SMTP server addresses exist. Define the objects using **Objects** > **Object Management**. You can also create these objects when you edit the policy.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **SMTP Server**.

Step 3 Select the network objects that identify the **Primary Server IP Address** and optionally, the **Secondary Server IP Address**.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

After completing this task, email alerts can be sent through Syslog using the specified SMTP servers.

Configure SNMP for a security device

Configure SNMP monitoring to allow network management stations to monitor the health and status of Cisco security devices. This task enables centralized visibility and management of devices using SNMP.

- Enable SNMP so that SNMP management stations can monitor devices.
- SNMP versions 1, 2c, and 3, including traps and read access, are supported. SNMP write access is not supported.

Simple Network Management Protocol (SNMP) defines a standard way for network management stations running on PCs or workstations to monitor the health and status of many types of devices, including switches, routers, and security appliances. You can use the SNMP page to configure a firewall device for monitoring by SNMP management stations.

SNMPv3 supports read-only users and encryption with DES(deprecated), 3DES, AES256, AES192, and AES128.

The DES option has been deprecated. If your deployment includes SNMP v3 users using DES encryption and that were created using a version earlier than 6.5, you can continue to use those users for Firewall Threat Defense devices running versions 6.6 or earlier. However, you cannot edit those users and retain DES encryption or create new users with DES encryption. If your Cloud-Delivered Firewall Management Center manages any Firewall Threat Defense devices running Versions 7.0+, deploying a platform settings policy that uses DES encryption to those Firewall Threat Defense devices will fail.



Note

- You can configure SNMP only on Routed and Diagnostic interfaces.
- To create an alert to an external SNMP server, access **Policies** > **Actions** > **Alerts**.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **SNMP**.

Step 3 Enable SNMP, and configure basic options.

- **Enable SNMP Servers**—Whether to provide SNMP information to the configured SNMP hosts. You can deselect this option to disable SNMP monitoring while retaining the configuration information.
- **Read Community String**, **Confirm**—Enter the password used by a SNMP management station when sending requests to the Firewall Threat Defense device. The SNMP community string (also known as password) is a shared secret among the SNMP management stations and the network nodes. The security device uses this password to determine if the incoming SNMP request is valid. The password is a case-sensitive alphanumeric string of up to 32 characters; spaces and special characters are not permitted.
- **System Administrator Name**—Enter the name of the device administrator or other contact person. This string is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
- **Location**—Enter the location of this security device (for example, Building 42, Sector 54). This string is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
- **Port**—Enter the UDP port on which incoming requests will be accepted. The default is 161.

Step 4 (SNMPv3 only.) Add SNMPv3 users, on page 1271.

Step 5 Add SNMP hosts, on page 1274.

Step 6 Configure SNMP traps, on page 1276.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

SNMP protocol

SNMP protocol is an application-layer protocol that

- facilitates the exchange of management information between network devices, and
- is part of the TCP/IP protocol suite.

SNMP protocol support and access

Firewall Threat Defense provides support for network monitoring using SNMP Versions 1, 2c, and 3, and supports the use of all three versions simultaneously. The SNMP agent runs on the Firewall Threat Defense interface. You can monitor network devices using management systems (NMSs) such as OpenView network management system. Firewall Threat Defense supports SNMP read-only access with a GET request. SNMP does not support write access or the SNMP SET request, so you cannot make changes with SNMP.

You can configure the Firewall Threat Defense to send traps, which are messages sent to the management station when certain events occur. These messages are called event notifications. You can also use the NMS to browse Management Information Bases (MIBs) on the security devices. MIBs are collections of definitions, and the Firewall Threat Defense maintains a database of values for each definition. To browse a MIB, issue a series of GET-NEXT or GET-BULK requests from the NMS to the MIB tree to get values.

The SNMP agent sends a notification to your management station when a specified event occurs, such as a network link going up or down. The notification includes an SNMP OID, which identifies the source device. The agent also replies when your management station requests information.

List SNMP terminology

This reference lists and defines the terms that are commonly used when working with SNMP.

Table 76: SNMP terminology

Term	Description
Agent	The SNMP server running on the Secure Firewall Threat Defense. The SNMP agent has these features: • responds to requests for information and actions from the network management station. • ontrols access to its Management Information Base, the collection of objects that the SNMP manager can view or change. • does not allow SET operations.
Browsing	Monitoring the health of a device from the network management station by polling required information from the SNMP agent on the device. This activity may include issuing a series of GET-NEXT or GET-BULK requests of the MIB tree from the network management station to determine values.
Management Information Bases (MIBs)	Standardized data structures for collecting information about packets, connections, buffers, failovers, and other items. MIBs are defined by the product, protocols, and hardware standards used by most network devices. SNMP network management stations can browse MIBs and request specific data or events be sent as they occur.
Network management stations (NMSs)	The PCs or workstations SET up to monitor SNMP events and manage devices.

Term	Description
Object identifier (OID)	The system that identifies a device to its NMS and indicates to users the source of information monitored and displayed.
Trap	Predefined events that generate a message from the SNMP agent to the NMS. Events include alarm conditions such as link-up, link-down, cold-start, warm-start, authentication events, or syslog messages. **coldStart**—The coldStart trap occurs when the SNMP agent starts after the SNMP configuration. This trap also occurs when the agent starts after a system reboot. **Note** For cluster and HA nodes, after a reload, if the interface reboot time exceeds 5 minutes (preset threshold), the trap is dropped. When the cluster and HA nodes have rebooted successfully, all other traps are sent as expected. The **SNMP-server enable traps SNMP coldstart** command is used to enable and disable transmission of these traps.

Describe MIBs and traps

MIBs are either standard or enterprise-specific. Standard MIBs are created by the IETF and described in various RFCs. Traps report significant events, such as errors or failures, on a network device. SNMP traps can be defined in standard or enterprise-specific MIBs. The IETF also creates standard traps, documented in various RFCs. SNMP traps are compiled into ASA software.

If needed, you can also download RFCs, standard MIBs, and standard traps from these locations:

http://www.IETF.org/

Browse the SNMP Object Navigator to look up Cisco MIBs, traps, and OIDs from this location:

https://snmp.cloudapps.cisco.com/Support/SNMP/do/BrowseOID.do?local=en

In addition, download Cisco OIDs by FTP from this location:

ftp://FTP.cisco.com/pub/mibs/oid/oid.tar.gz

List of supported tables and objects in MIBs

This reference lists the supported tables and objects for the specified MIBs, including details for each MIB such as counters, OIDs, and descriptions.

Remote Access VPN polling

Table 77: CISCO-REMOTE-ACCESS-MONITOR-MIB

Counter	OID	Description
Active Sessions	crasNumSessions (1.3.6.1.4.1.9.9.392.1.3.1)	The number of currently active sessions.
Users	crasNumUsers (1.3.6.1.4.1.9.9.392.1.3.3)	The number of users who have active sessions.

Counter	OID	Description
Peak Sessions	crasNumPeakSessions (1.3.6.1.4.1.9.9.392.1.3.41)	The number of peak RA sessions since system up.

Site-to-site VPN tunnel polling

Table 78: CISCO-REMOTE-ACCESS-MONITOR-MIB

Counter	OID	Description
LAN to LAN Sessions	crasL2LNumSessions (1.3.6.1.4.1.9.9.392.1.3.29)	The number of currently active LAN to LAN sessions.
Peak LAN to LAN Sessions	crasL2LPeakConcurrentSessions (1.3.6.1.4.1.9.9.392.1.3.31)	The number of peak concurrent LAN to LAN sessions since the system is up.

Connection polling

Table 79: CISCO-FIREWALL-MIB

Counter	OID	Description
Active Connections	cfwConnectionActive (1.3.6.1.4.1.9.9.147.1.2.2.2.1.3.40.6)	The number of connections currently in use by the entire firewall.
Peak Connections	cfwConnectionPeak (1.3.6.1.4.1.9.9.147.1.2.2.2.1.3.40.7)	The highest number of connections in use at any one time since system startup.
Connections Per Second	cfwConnectionPerSecond (1.3.6.1.4.1.9.9.147.1.2.2.3)	The current connections per second rate on the firewall.
Peak Connections Per Second	cfwConnectionPerSecondPeak (1.3.6.1.4.1.9.9.147.1.2.2.4)	The highest number of connections per second on the firewall since system startup.

NAT translation polling

Table 80: CISCO-NAT-EXT-MIB

Counter	OID	Description
Active Translations	cneAddrTranslationNumActive (1.3.6.1.4.1.9.9.532.1.1.1.1)	The total number of address translation entries that are currently available in the NAT device. This indicates the aggregate of the translation entries created from both the static and dynamic address translation mechanisms.
Peak Active Translations	cneAddrTranslationNumPeak (1.3.6.1.4.1.9.9.532.1.1.1.2)	The maximum number of address translation entries that are active at any one time since system startup. This value indicates the high watermark of address translation entries active at any one time since system startup. This object includes the translation entries created from both the static and dynamic address translation mechanisms.

Routing table entries polling

Table 81: IP-FORWARD-MIB

Counter	OID	Description
Active Translations	inetCidrRouteNumber (1.3.6.1.2.1.4.24.6)	The total number of current inetCidrRouteTable entries that are valid.

Interface duplex status polling

Table 82: CISCO-IF-EXTENSION-MIB

Counter	OID	Description
Duplex Status	cieIfDuplexCfgStatus (1.3.6.1.4.1.9.9.276.1.1.2.1.20)	This object specifies the configured duplex status for the given interface.
Detected Duplex Status	cieIfDuplexDetectStatus (1.3.6.1.4.1.9.9.276.1.1.2.1.21)	This object specifies the detected duplex status for the given interface.

Snort 3 intrusion EVENT rate polling

Table 83: CISCO-UNIFIED-FIREWALL-MIB

Counter	OID	Description
Snort 3 Intrusion event Rate	cufwAaicIntrusionEvtRate (1.3.6.1.4.1.9.9.491.1.5.3.2.1)	The rate at which intrusion events were recorded by Snort on this firewall, averaged over the last 300 seconds.

BGP peer-flap trap notification

Table 84: BGP4-MIB

Counter	OID	Description
BGP Peer-flap	bgpBackwardTransition (1.3.6.1.4.1.9.9.491.1.5.3.2.1)	The BGPBackwardTransition event is generated when the BGP FSM moves from a higher numbered state to a lower numbered state.

CPU utilization polling

Table 85: CISCO-PROCESS-MIB

Counter	OID	Description
CPU Total Utilization	cpmCPUTotal1minRev (1.3.6.1.4.1.9.9.109.1.1.1.1.7.1)	Total system CPU utilization for the last one minute.

Counter	OID	Description
Individual CPU Core Utilization	Associated parameters and values of cpmCPUTotal1minRev 1.3.6.1.4.1.9.9.109.1.1.1.1.7.2 to 1.3.6.1.4.1.9.9.109.1.1.1.1.7.(n+1)	Individual CPU core utilization values for the last one minute, where 'n' represents the number of cores. Examples: • 3.6.1.4.1.99.109.1.1.1.1.7.(n+2) - Aggregate system CPU utilization % (This value is the same as the system CPU usage from .3.6.1.4.1.99.109.1.1.1.1.7.1 in single context mode). • 3.6.1.4.1.99.109.1.1.1.1.7.(n+3) - Snort average CPU utilization % (total aggregate value of all snort instances) • 3.6.1.4.1.99.109.1.1.1.1.7.(n+4) - System PROCESS average % (average of "Sysproc" cores)

Note The SNMP OIDs 1.3.6.1.2.1.25.3.3 and 1.3.6.1.2.1.25.3.4 pertaining to CPU monitoring (hrProcessorTable and hrNetworkTable) were removed on ASA FirePOWER. You can view and MONITOR the CPU health details of the device only through its device manager.

In ENTITY-MIB, two new vendor OIDs were added for the dual EPM 2X100G and 4X200G cards of Secure Firewall 4200 for slot 2 and slot 3—cevFPRNM4X200Gng and cevFPRNM2X100Gng.

Add SNMPv3 users

Add SNMPv3 users to provide secure, authenticated, and encrypted SNMP management access to your device. This task ensures that only authorized users can access SNMPv3 features using specified authentication and encryption algorithms.

You create users for SNMPv3 only. These steps are not applicable for SNMPv1 or SNMPv2c.

Note SNMPv3 only supports read-only users.

SNMP users have a specified username, an authentication password, an encryption password, and authentication and encryption algorithms to use.

Note If you add a new cluster unit or replace a High Availability unit when using SNMPv3 with clustering or High Availability, SNMPv3 users are not replicated to the new unit. Remove the users, add them again, and redeploy your configuration to replicate users to the new unit.

For a successful SNMPv3 user authentication, the authentication key size should be equal to or larger than the encryption key size. For example, a user authentication using SHA-1 with AES256 fails, whereas authentication using SHA256 with AES256 succeeds.

The authentication algorithm options are MD5 (deprecated, pre-6.5 only), SHA, SHA224, SHA256, and SHA384.

Note The MD5 option has been deprecated. If your deployment includes SNMP v3 users using the MD5 authentication algorithm that were created using a version previous to 6.5, you can continue to use those users for FTD devices running versions 6.7 and previous. However, you cannot edit those users and retain the MD5 authentication algorithm, or create new users with the MD5 authentication algorithm. If your management center manages any threat defenses running Versions 7.0+, deploying a platform settings policy that uses the MD5 authentication algorithm to those threat defenses will fail.

The encryption algorithm options are DES (deprecated, pre-6.5 only), 3DES, AES256, AES192, and AES128.

Note The DES option has been deprecated. If your deployment includes SNMP v3 users using DES encryption that were created using a version previous to 6.5, you can continue to use those users for threat defenses running versions 6.7 and previous. However, you cannot edit those users and retain DES encryption, or create new users with DES encryption. If your management center manages any threat defenses running Versions 7.0+, deploying a platform settings policy that uses DES encryption to those threat defenses will fail.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Choose **SNMP** > **Users** tab.

Step 3 Click **Add** and completed the following:

a) Select the security level for the user from the **Security Level** drop-down list.

- **Auth**—Authentication but No Privacy, which means that messages are authenticated.
- **No Auth**—No Authentication and No Privacy, which means that no security is applied to messages.
- **Priv**—Authentication and Privacy, which means that messages are authenticated and encrypted.

b) Enter the name of the SNMP user in the **Username** field. Usernames must be 32 characters or fewer.

c) From the **Encryption Password Type** drop-down list, choose one of these options:

- **Clear text**—The Firewall Threat Defense device encrypts the password when you deploy the policy. In High Availability or cluster deployments, choose this option so that each unit localizes the SNMPv3 key with its own SNMP engineID.
- **Encrypted**—The Firewall Threat Defense device directly deploys the precomputed localized key in hexadecimal format.

Note
Do not choose **Encrypted** for SNMPv3 users in High Availability or cluster deployments. An encrypted SNMPv3 password is a localized key that is generated from a specific SNMP engineID. Each High Availability or cluster unit has a unique SNMP engineID. If you deploy a localized key generated for one unit to its peer, SNMPv3 polling to the peer fails with an authentication error.

Choose **Clear text** so that each unit localizes the key during deployment.

d) In the **Auth Algorithm Type** drop-down list, select the type of authentication you want to use: SHA, SHA224, SHA256, or SHA384.

Note
The MD5 option has been deprecated. If your deployment includes SNMP v3 users using the MD5 authentication algorithm that were created using a version previous to 6.5, you can continue to use those users for FTD devices running versions 6.7 and previous. However, you cannot edit those users and retain the MD5 authentication algorithm, or create new users with the MD5 authentication algorithm. If your Cloud-Delivered Firewall Management Center manages any Firewall Threat Defenses running Versions 7.0+, deploying a platform settings policy that uses the MD5 authentication algorithm to those Firewall Threat Defenses will fail.

e) In the **Authentication Password** field, enter the password to use for authentication. If you selected Encrypted as the Encrypt Password Type, the password must be formatted as xx:xx:xx..., where xx are hexadecimal values.

Note
The length of the password will depend on the authentication algorithm selected. For all passwords, the length must be 256 characters or less.

If you selected Clear Text as the Encrypt Password Type, repeat the password in the **Confirm** field.

f) In the **Encryption Type** drop-down list, select the type of encryption you want to use: AES128, AES192, AES256, 3DES.

Note
To use AES or 3DES encryption, you must have the appropriate license installed on the device.

Note
The DES option has been deprecated. If your deployment includes SNMP v3 users using DES encryption that were created using a version previous to 6.5, you can continue to use those users for Firewall Threat Defenses running versions 6.7 and previous. However, you cannot edit those users and retain DES encryption, or create new users with DES encryption. If your Cloud-Delivered Firewall Management Center manages any Firewall Threat Defenses running Versions 7.0+, deploying a platform settings policy that uses DES encryption to those Firewall Threat Defenses will fail.

g) Enter the password to use for encryption in the **Encryption Password** field. If you selected Encrypted as the Encrypt Password Type, the password must be formatted as xx:xx:xx..., where xx are hexadecimal values. For encrypted passwords, the length of the password depends on the encryption type selected. The password sizes are as follows. Each xx is one octal.:

- AES 128 requires 16 octals
- AES 192 requires 24 octals
- AES 256 requires 32 octals
- 3DES requires 32 octals
- DES can be any size

Note
For all passwords, the length must be 256 characters or fewer.

If you selected Clear Text as the Encrypt Password Type, repeat the password in the **Confirm** field.

Step 4 Click **OK**.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

After completing these steps, SNMPv3 users are added with the specified authentication and encryption settings. The device is now configured for secure SNMPv3 management access.

Add SNMP hosts

Define SNMP hosts and their communication parameters to enable SNMP management stations to access and monitor the device. This task allows you to specify which hosts are permitted and how they communicate with the device.

Use Host to add or edit entries in the SNMP Hosts table on the SNMP page. These entries represent SNMP management stations allowed to access the Firewall Threat Defense device.

You can add up to 8192 hosts. Only 128 of this number can be for traps.

Note In 7.4 and later, the Management and Diagnostic interfaces are merged. If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, then you must use separate Platform Settings policies for merged and non-merged devices (7.3 and earlier, and some upgraded 7.4 FTDs).

Before you begin

Ensure that the network objects that define the SNMP management stations exist. Select **Device** > **Object Management** to configure network objects.

Note The supported network objects include IPv6 hosts, IPv4 hosts, IPv4 range and IPv4 subnet addresses.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Choose **SNMP** > **Hosts** tab.

Step 3 Click **Add** and completed the following:

a) In the **IP Address** field, either enter a valid IPv6 or IPv4 host or select the network object that defines the SNMP management station's host address.

The IP address can be an IPv6 host, IPv4 host, IPv4 range or IPv4 subnet.

b) Select the appropriate SNMP version from the **SNMP version** drop-down list.

c) (SNMPv3 only.) Select the username of the SNMP user that you configured from the **User Name** drop-down list.

Note
You can associate up to 23 SNMP users per SNMP host.

d) (SNMPv1, 2c only.) In the **Read Community String** field, enter the community string that you have already configured, for read access to the device. Re-enter the string to confirm it.

Note
This string is required, only if the string used with this SNMP station is different from the one already defined in the **Enable SNMP Server** section.

e) Select the type of communication between the device and the SNMP management station. You can select both types. Choose from:

- **Poll**—The management station periodically requests information from the device.
- **Trap**—The device sends trap events to the management station as they occur.

Note
When the SNMP host IP address is either an IPv4 range or an IPv4 subnet, you can configure either **Poll** or **Trap**, not both.

f) In the **Port** field, enter a UDP port number for the SNMP host. The default value is 162. The valid range is 1 to 65535.

g) Select the interface type for communication between the device and the SNMP management station under the **Reachable By** options. You can select either the Management interface of the device or an available security zone/named interface. Choose from:

- **Device Management Interface**—Communication between the device and the SNMP management station occurs over the Management interface.
  - When you choose this interface for SNMPv3 polling, all configured SNMPv3 users are allowed to poll and are not restricted to the user chosen in 3.c, on page 1275. Here, SNMPv1 and SNMPv2c are not allowed from an SNMPv3 host.

- When you choose this interface for SNMPv1 and SNMPv2c polling, the polling is not restricted at all to the version selected in 3.b, on page 1275.

- **Security Zones or Named Interface**—Communication between the device and the SNMP management station occurs over a security zone or interface.
  - Search for zones in the **Available Zones** field.
  - Add the zones that contain the interfaces through which the device communicates with the management station to the **Selected Zone/Interface** field. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zone/Interface** list and click **Add**. You can also choose a loopback interface and virtual-router-aware interfaces. The host will be configured on a device only if the device includes the selected interfaces or zones.

Note
VTI tunnel interfaces are not supported for SNMP.

Step 4 Click **OK**.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure SNMP traps

Use SNMP traps to configure event notifications for Firewall Threat Defense. Traps are unsolicited notifications sent from the device to the management station for certain events, such as linkup, linkdown, and syslog event generation. Each SNMP trap includes an object ID (OID) for your device.

Some traps are not applicable to certain hardware models. These traps will be ignored if you apply the policy to one of these models. For example, not all models have field-replaceable units, so the **Field Replaceable Unit Insert/Delete** trap will not be configured on those models.

SNMP traps are defined in standard MIBs or in enterprise-specific MIBs. Standard traps are created by the IETF and documented in various RFCs. The Firewall Threat Defense software includes SNMP traps.

If needed, you can download RFCs, standard MIBs, and standard traps from http://www.IETF.org/.

Browse the complete list of CISCO MIBs, traps, and OIDs using the SNMP Object Navigator on the Cisco website: SNMP Object Navigator.

In addition, download CISCO OIDs by FTP: ftp://ftp.cisco.com/pub/mibs/oid/oid.tar.gz.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **SNMP** > **SNMP Traps** to configure SNMP traps (event notifications) for the Firewall Threat Defense device.

Step 3 Select the appropriate Enable Traps options. You can select either or both options.

a) Check **Enable All SNMP Traps** to select all traps in the subsequent four sections.
b) Check **Enable All Syslog Traps** to enable transmission of trap-related syslog messages.

Note
SNMP traps are of higher priority than other notification messages because they are expected to be near real-time. When you enable all SNMP or syslog traps, it is possible for the SNMP process to consume excess resources in the agent and in the network, causing the system to hang. If you notice system delays, unfinished requests, or timeouts, you can selectively enable SNMP and syslog traps. You can also limit the rate at which syslog messages are generated by severity level or message ID. For example, all syslog message IDs that begin with the digits 212 are associated with the SNMP class; see Configure syslog rate limits, on page 1296.

Step 4 The event-notification traps in the **Standard** section are enabled by default for an existing policy:

- **Authentication** – Unauthorized SNMP access. This authentication failure occurs for packets with an incorrect community string.
- **Link Up** – One of the device's communication links has become available (it has "come up"), as indicated in the notification.
- **Link Down** – One of the device's communication links has failed, as indicated in the notification.
- **Cold Start** – The device is reinitializing itself such that its configuration or the protocol entity implementation may be altered.
- **Warm Start** – The device is reinitializing itself such that its configuration and the protocol entity implementation is unaltered.

Step 5 Select the desired event-notification traps in the **Entity MIB** section:

- **Field Replaceable Unit Insert** – A Field Replaceable Unit (FRU) has been inserted, as indicated. (FRUs include assemblies such as power supplies, fans, processor modules, interface modules, etc.)
- **Field Replaceable Unit Delete** – A Field Replaceable Unit (FRU) has been removed, as indicated in the notification
- **Configuration Change** – There has been a hardware change, as indicated in the notification

Step 6 Select the desired event-notification traps in the **Resource** section:

- **Connection Limit Reached** – This trap indicates that a connection attempt was rejected because the configured connections limit has been reached.

Step 7 Select the desired event-notification traps in the **Other** section:

- **NAT Packet Discard** – This notification is generated when IP packets are discarded by the NAT function. Available Network Address Translation addresses or ports have fallen below configured threshold.
- **CPU Rising Threshold** – This notification is generated when rising CPU utilization exceeds a predefined threshold for a configured period of time. Check this option to enable CPU rising threshold notifications:
  - **Percentage** – The default value is 70% for the high threshold notification; the range is 10 to 94%. The critical threshold is 95%.
  - **Period** – The default monitoring period is 1 minute; the range is 1 to 60 minutes.

- **Memory Rising Threshold** – This notification is generated when rising memory utilization exceeds a predefined threshold, thus reducing available memory. Check this option to enable memory rising threshold notifications:
  - **Percentage** – The default value is 70% for the high threshold notification; the range is 50 to 95%.
- **Failover** – This notification is generated when there is a change in the failover state as reported by the CISCO-UNIFIED-FIREWALL-MIB.
- **Cluster** – This notification is generated when there is a change in the cluster health as reported by the CISCO-UNIFIED-FIREWALL-MIB.
- **Peer Flap** – This notification is generated when there is BGP route flapping, a situation in which BGP systems send an excessive number of update messages to advertise network reachability information.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

After completing this task, SNMP traps are configured on the Firewall Threat Defense device, and event notifications will be sent to the management station for the selected events. This enables proactive monitoring and timely response to device events.

Configure SSL settings

This task enables you to configure SSL settings for remote access VPN sessions, allowing you to specify protocol versions, security levels, and available algorithms to ensure secure communication.

You must have administrator privileges and be in a leaf domain to perform this task.

Make sure that you are running a fully licensed version of the Cloud-Delivered Firewall Management Center. The SSL Settings are disabled if you are running Cloud-Delivered Firewall Management Center in evaluation mode. SSL Settings are also disabled if the licensed Cloud-Delivered Firewall Management Center version does not meet the export-compliance criteria. If you are using Remote Access VPN with SSL, your Smart Account must have the strong-crypto features enabled. For more information, see *License Types and Restrictions* in the *Cisco Secure Firewall Management Center Administration Guide*.

Before you begin

Ensure you have administrator privileges and are in a leaf domain. Verify that you are running a fully licensed version of the management center with strong-crypto features enabled if using Remote Access VPN with SSL.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit a Firewall Threat Defense policy.

Step 2 Select **SSL**.

Step 3 Add entries to the **Add SSL Configuration** table.

a) Click **Add** to create a new entry, or click **Edit** if the entry already exists.

b) Select the required security configurations from the drop-down list.

- **Protocol Version**—Specifies the TLS protocols to be used while establishing remote access VPN sessions.
- **Security Level**—Indicates the kind of security positioning you would like to set up for the SSL.

Step 4 Select the **Available Algorithms** based on the protocol version that you select and click **Add** to include them for the selected protocol. For more information, see Configure SSL settings, on page 1279.

The algorithms are listed based on the protocol version that you select. Each security protocol identifies a unique algorithm for setting up the security level.

Step 5 Click **OK** to save the changes.

What to do next

Select **Deploy** > **Deployment** and click **Deploy** to deploy the policy to the assigned devices.

Configure SSL settings

The Firewall Threat Defense device uses the Secure Sockets Layer (SSL) protocol and Transport Layer Security (TLS) to support secure message transmission for Remote Access VPN connection from remote clients. The SSL Settings window lets you configure SSL versions and encryption algorithms that will be negotiated and used for message transmission during remote VPN access over SSL.

Note Though you configure Cloud-Delivered Firewall Management Center and Firewall Threat Defense to operate in a security certifications (UCAPL, CC, or FIPS) compliance mode, the Cloud-Delivered Firewall Management Center allows configuration of unsupported ciphers. For example, in a FIPS enabled mode, management center allows configuring DH Group 5 which is not FIPS-compliant. However, VPN tunnel does not negotiate due to the non-compliant cipher usage.

Configure the SSL Settings at this location:

Devices > **Platform Settings**, and then click **SSL**.

Fields

Minimum SSL Version as Server—Specify the minimum SSL/TLS protocol version that the Firewall Threat Defense device uses when acting as a server. For example, when it functions as a Remote Access VPN Gateway.

TLS Version—Select one of these TLS versions from the drop-down list:

TLS V1	Accepts SSLv2 client hellos and negotiates TLSv1 (or greater).
TLSV1.1	Accepts SSLv2 client hellos and negotiates TLSv1.1 (or greater).
TLSV1.2	Accepts SSLv2 client hellos and negotiates TLSv1.2 (or greater).
TLSV1.3	Accepts SSLv2 client hellos and negotiates TLSv1.3 (or greater).

Note TLS 1.3 in remote access VPN requires Cisco Secure Client, Version 5.0 and above.

DTLS Version—Select the DTLS versions from the drop-down list, based on the selected TLS version. By default, DTLSv1 is configured on Firewall Threat Defense devices, you can choose the DTLS version as per your requirement.

Note Ensure that the TLS protocol version is higher than or equal to the DTLS protocol version selected. TLS protocol versions support these DTLS versions:

TLS V1	DTLSv1
TLSV1.1	DTLSv1
TLSV1.2	DTLSv1, DTLSv1.2
TLSV1.3	DTLSv1, DTLSv1.2

Diffie-Hellman Group—Choose a group from the drop-down list. Available options are Group1 - 768-bit modulus, Group2 - 1024-bit modulus, Group5 - 1536-bit modulus, Group14 - 2048-bit modulus, 224-bit prime order, and Group24 - 2048-bit modulus, 256-bit prime order. The default is Group1.

Elliptical Curve Diffie-Hellman Group—Choose a group from the drop-down list. Available options are Group19 - 256-bit EC, Group20 - 384-bit EC, and Group21 - 521-bit EC. The default value is Group19.

TLSv1.2 adds support for these ciphers:

- ECDHE-ECDSA-AES256-GCM-SHA384
- ECDHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES256-GCM-SHA384
- AES256-GCM-SHA384
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES256-SHA384
- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-RSA-AES128-GCM-SHA256
- DHE-RSA-AES128-GCM-SHA256
- RSA-AES128-GCM-SHA256

- ECDHE-ECDSA-AES128-SHA256
- ECDHE-RSA-AES128-SHA256

Note
ECDSA and DHE ciphers are the highest priority.

TLSv1.3 adds support for these ciphers:

- TLS_AES_128_GCM_SHA256
- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_256_GCM_SHA384

The SSL configuration table can be used to specify the protocol version, security level, and Cipher algorithms that you want to support on the Secure Firewall Threat Defense devices.

Protocol Version—Lists the protocol version that the Secure Firewall Threat Defense device supports and uses for SSL connections. Available protocol versions are:

- Default
- TLSV1
- TLSV1.1
- TLSV1.2
- TLSV1.3
- DTLSv1
- DTLSv1.2

Security Level—Lists the cipher security levels that Firewall Threat Defense device supports and uses for SSL connections.

If you have Firewall Threat Defense devices with evaluation license, the security level is Low by default. With Firewall Threat Defense smart license, the default security level is High. You can choose one of these options to configure the required security level:

- **All** includes all ciphers, including NULL-SHA.
- **Low** includes all ciphers, except NULL-SHA.
- **Medium** includes all ciphers, except NULL-SHA, DES-CBC-SHA, RC4-SHA, and RC4-MD5 (this is the default).
- **FIPS** includes all FIPS-compliant ciphers, except NULL-SHA, DES-CBC-SHA, RC4-MD5, RC4-SHA, and DES-CBC3-SHA, TLS_CHACHA20_POLY1305_SHA256.
- **High** includes only AES-256 with SHA-2 ciphers and applies to TLS version 1.2 and the *default* version.
- **Custom** includes one or more ciphers that you specify in the Cipher algorithms/custom string box. This option provides you with full control of the cipher suite using OpenSSL cipher definition strings.

Cipher Algorithms/Custom String—Lists the cipher algorithms that the Firewall Threat Defense device supports and uses for SSL connections. For more information about ciphers using OpenSSL, see https://www.openssl.org/docs/apps/ciphers.html

Ciphers supported by TLSv1.2 only

- ECDHE-ECDSA-AES256-GCM-SHA384
- ECDHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES256-GCM-SHA384
- AES256-GCM-SHA384
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES256-SHA384
- DHE-RSA-AES256-SHA256
- AES256-SHA256
- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-RSA-AES128-GCM-SHA256
- DHE-RSA-AES128-GCM-SHA256
- AES128-GCM-SHA256
- ECDHE-ECDSA-AES128-SHA256
- ECDHE-RSA-AES128-SHA256
- DHE-RSA-AES128-SHA256
- AES128-SHA256



Note ECDSA and DHE ciphers are the highest priority.

Ciphers not supported by TLSv1.1 or TLSv1.2

- RC4-SHA
- RC4-MD5

- DES-CBC-SHA
- NULL-SHA

Syslog

Logging information

Logging information can help you identify and isolate network or device configuration problems.

- System logging (syslog) can be enabled for Firewall Threat Defense devices.
- Some security events can be sent to a syslog server.

Syslog messages

A syslog message is a log event that

- collects messages from devices to a server running a syslog daemon
- enables aggregation of logs and alerts for centralized management, and
- provides protected long-term storage for logs, supporting troubleshooting and incident handling.

System log configuration details

Devices can send their log messages to a UNIX-style syslog service. A syslog service accepts messages and either stores them in files or prints them, based on a simple configuration file.

Table 86: System logs for Secure Firewall Threat Defense

Logs related to	Details	Configure in
Device and system health, network configuration	This syslog configuration generates messages for features running on the data plane, that is, features that are defined in the CLI configuration that you can view with the **show running-config** command. This includes features such as routing, VPN, data interfaces, DHCP server, NAT, and so forth. Data plane syslog messages are numbered, and they are the same as those generated by devices running ASA software. However, Secure Firewall Threat Defense does not necessarily generate every message type that is available for ASA Software. For information on these messages, see *Cisco Secure Firewall Threat Defense syslog messages* at https://www.cisco.com/c/en/us/td/docs/security/firepower/Syslogs/b_fptd_syslog_guide.html. This configuration is explained in these topics.	**Platform settings**

Logs related to	Details	Configure in
Security events	This syslog configuration generates alerts for file and malware, connection, Security Intelligence, and intrusion events.	**Platform settings** and the **Logging** in an access control policy
(All devices) Policies, rules, and events	This syslog configuration generates alerts for access control rules, intrusion rules, and other advanced services as described in *Configurations supporting alert responses* in the *Cisco Secure Firewall Management Center Administration Guide*. These messages are not numbered. For information on configuring this type of syslog, see *Creating a syslog alert response* in the *Cisco Secure Firewall Management Center Administration Guide*.	**alert responses** and the **Logging** in an access control policy

You can set up multiple syslog servers and control the messages and events sent to each server. You can also configure varied destinations, such as console, email, or internal buffer.

Syslog message severity levels

The syslog message severity levels define the importance of system messages.

Table 87: Syslog message severity levels

Level Number	Severity Level	Description
0	**emergencies**	You cannot use the system.
1	**alert**	You need to take immediate action.
2	**critical**	The system has critical conditions to be addressed.
3	**error**	Error conditions.
4	**warning**	The system has warnings you should review.
5	**notification**	The system reports significant events you should monitor.
6	**informational**	Informational messages only.
7	**debugging**	The system generates only debugging messages Log at this level only temporarily, when debugging issues. This log level can generate many messages, which might affect system performance.

Note Firewall Threat Defense does not generate syslog messages with a severity level of zero (emergencies).

Syslog message filters

Syslog message filter criteria

You can direct syslog messages to an output destination according to these criteria:

- Syslog message ID number

  (This does not apply to syslog messages for security events such as connection and intrusion events.)

- Syslog message severity level
- Syslog message class (equivalent to a functional area)

  (This rule does not apply to syslog messages for security events, including connection and intrusion events.)

You can create a message list with specific criteria, which you can then specify as the output destination is set. Alternatively, you can configure the Firewall Threat Defense device to send a particular message class to each type of output destination independently of the message list.

(Message lists do not apply to syslog messages for security events such as connection and intrusion events.)

Syslog message classes and associated message ID numbers



Note This topic does not apply to messages for security events (connection, intrusion, etc.)

You can use syslog message classes in two ways:

- Specify an output location for an entire category of syslog messages.
- Create a message list that specifies the message class.

Use the syslog message class to categorize syslog messages by type, such as features or functions of the device. For example, use the RIP class for RIP routing.

Syslog messages in the same class have the same three initial digits in their message ID numbers. For example, the syslog messages for the VPN client feature use ID numbers 611101 to 611323.

Most ISAKMP syslog messages include prepended objects that identify the tunnel. When available, these objects appear before the descriptive text of a syslog message. If the object is unknown when the syslog message is generated, the heading and value combination is omitted.

The objects are prefixed in this manner:

Group = *groupname*, Username = *user*, IP = *IP_address*

The term 'group' here refers to the tunnel-group configuration. The username is from the local database or AAA server. The IP address is the public IP address of the remote access client or Layer 2 peer.

See the message classes and the range of message ID numbers in each class in this table.

Table 88: Syslog message classes and associated message ID numbers

Class	Definition	Syslog Message ID Numbers
auth	User Authentication	109, 113
—	Access Lists	106
—	Application Firewall	415
—	Botnet Traffic Filtering	338
bridge	Transparent Firewall	110, 220
ca	PKI Certification Authority	717
citrix	Citrix Client	723
—	Clustering	747
—	Card Management	323
config	Command Interface	111, 112, 208, 308
csd	Secure Desktop	724
cts	Cisco TrustSec	776
dap	Dynamic Access Policies	734
EAP, eapoudp	EAP or EAPoUDP for Network Admission Control	333, 334
EIGRP	EIGRP Routing	336
email	E-mail Proxy	719
—	Environment Monitoring	735
ha	Failover	101, 102, 103, 104, 105, 210, 311, 709
—	Identity-based Firewall	746
ids	Intrusion Detection System	400, 733
—	IKEv2 Toolkit	750, 751, 752
IP	IP Stack	209, 215, 313, 317, 408
ipaa	IP Address Assignment	735
ips	Intrusion Protection System	400, 401, 420
—	IPv6	325
—	Licensing	444

Class	Definition	Syslog Message ID Numbers
MDM-proxy	MDM Proxy	802
NAC	Network Admission Control	731, 732
nacpolicy	NAC Policy	731
nacsettings	NAC Settings to apply NAC Policy	732
—	NAT and PAT	305
—	Network Access Point	713
NP	Network Processor	319
—	NP SSL	725
OSPF	OSPF Routing	318, 409, 503, 613
—	Password Encryption	742
—	Phone Proxy	337
RIP	RIP Routing	107, 312
rm	Resource Manager	321
—	Smart Call Home	120
session	User Session	106, 108, 201, 202, 204, 302, 303, 304, 305, 314, 405, 406, 407, 500, 502, 607, 608, 609, 616, 620, 703, 710
SNMP	SNMP	212
—	ScanSafe	775
SSL	SSL Stack	725
svc	SSL VPN Client	722
sys	System	199, 211, 214, 216, 306, 307, 315, 414, 604, 605, 606, 610, 612, 614, 615,701, 711, 741
—	Threat Detection	733
tag-switching	Service Tag Switching	779
transactional-rule-engine-tre	Transactional Rule Engine	780
uc-IMS	UC-IMS	339
vm	VLAN Mapping	730

Class	Definition	Syslog Message ID Numbers
vpdn	PPTP and L2TP Sessions	213, 403, 603
VPN	IKE and IPsec	316, 320, 402, 404, 501, 602, 702, 713, 714, 715
vpnc	VPN Client	611
vpnfo	VPN Failover	720
vpnlb	VPN Load Balancing	718
—	VXLAN	778
webfo	WebVPN Failover	721
webvpn	WebVPN and Secure Client	716

Guidelines for logging

Restrictions for IPv6 logging

- IPv6 is supported. Syslog message can be sent using TCP or UDP.
- Ensure that the interface configured for sending syslogs is enabled, IPv6 capable, and the syslog server is reachable through the designated interface.
- You cannot use IPv6 for secure logging.

Syslog server configuration guidelines

- The syslog server must run a server program called syslogd. Windows includes a syslog server as part of its operating system.
- The syslog server operates based on the syslog-ng process of the firewall system. Do not use external configuration files, such as the *scwx.conf* file from SecureWorks. Because these files are not compatible with the device, using them causes parsing errors and the syslog-ng process fails.
- You can configure up to 16 syslog servers.
- Assign only one list or class to each syslog server and location.
- Do not configure Cloud-Delivered Firewall Management Center as a primary syslog server. The Cloud-Delivered Firewall Management Center can log some syslogs. However, it does not have adequate storage provision to accommodate large amounts of information from connection events for every sensor, especially when multiple sensors are used and all send syslogs.
- The syslog server should be reachable through the Firewall Threat Defense device. Configure the device to deny ICMP unreachable messages on the interface through which the syslog server is reachable, and to send syslogs to the same server. Enable logging for all severity levels. To prevent the syslog server from crashing, suppress the generation of syslogs 313001, 313004, and 313005.
- To view logs generated by the Firewall Threat Defense device, you must specify a logging output destination. If you enable logging without specifying a logging output destination, the Firewall Threat

Defense device generates messages but does not save them to a location from which you can view them. You must specify each different logging output destination separately.

- The Firewall Threat Defense syslog supports one-way TLS for secure logging. It does not support client certificate authentication, trustpoint configuration for syslog, or mutual TLS (X.509 certificate authentication).

Interfaces for syslog configuration

- Determining the egress interface for the syslog:
  - If the specified management-only interface has management-access enabled, the management center will perform route table lookups and determine the egress interface (could be data or management) based on best routing logic.
  - If you configure a management-only interface as logging host, that does not have management-access enabled, the management center will use the interface regardless of routing table entries.

    Thus, for the management center to always uses a dedicated management path for syslog traffic, configure the management interface without management-access, and then specify the interface in the logging host:

    ```
    interface <management-interface>
    management-only ----->Do not include management-access
    logging host <management-interface> <syslog-server-ip>
    ```

- When the syslog rate exceeds 50,000 messages per second, ensure that a data interface is used as egress interface rather than a management interface.
- Configure the data interface as the syslog source interface for production monitoring and auditing. Use the management interface only for temporary testing. This approach ensures optimal syslog reliability and prevents rate-limiting on the management interface.

TCP and UDP connections for syslog

- If you use TCP as the transport protocol, the system opens 4 connections to the syslog server to ensure that messages are not lost. If you collect messages from many devices and the combined connection overhead is too much for the server, use UDP instead.
- The number of UDP connections for syslog is directly related to the number of CPUs on the hardware platform and the number of syslog servers you configure. At any point in time, there can be as many UDP syslog connections as there are CPUs times the number of configured syslog servers. The global UDP connection idle timeout applies to these sessions, and the default duration is two minutes. You can adjust this setting to close sessions more quickly, but the timeout applies to all UDP connections, not just syslog. This is the intended behavior.
- When the Firewall Threat Defense device sends syslogs via TCP, the connection takes about one minute to initiate after the syslogd service restarts.
- When the TCP logging host goes down, it takes approximately six minutes to change its connection status from *Connected* to *Not connected*. Logging relies on TCP to detect the channel state; until then, logging sends the logs through the channel. During this time, when you execute the **show log**, the output would display the TCP logging host as connected. Once the TCP channel is closed, the TCP logging host state is updated to *Not connected*.

Configure syslog logging for Firewall Threat Defense devices

This task enables you to configure syslog logging for Firewall Threat Defense devices, allowing you to monitor and manage system events efficiently.

Syslog logging helps track security events and system activity on your devices. If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages.

See Firewall Threat Defense platform settings that apply to security event syslog messages , on page 1291 for more information about platform settings that apply to security event syslog messages.

Before you begin

See requirements in Guidelines for logging, on page 1288.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **Syslog** from the table of contents.

Step 3 Click **Logging Setup** to enable logging, specify FTP server settings, and specify Flash usage. For more information, see Configure logging and basic settings, on page 1291.

Step 4 Click **Logging Destinations** to enable logging to specific destinations and to specify filtering on message severity level, event class, or on a custom event list. For more information, see Enable logging destinations, on page 1293.

You must enable a logging destination to see messages at that destination.

Step 5 Click **E-mail Setup** to specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages. For more information, see Configure syslog messages to be sent to an e-mail address, on page 1294.

Step 6 Click **Events List** to define a custom event list that includes an event class, a severity level, and an event ID. For more information, see Create a custom event list, on page 1295.

Step 7 Click **Rate Limit** to specify the volume of messages being sent to all configured destinations and define the message severity level to which you want to assign rate limits. For more information, see Configure syslog rate limits, on page 1296.

Step 8 Click **Syslog Settings** to specify the logging facility, enable the inclusion of a time stamp, and enable other settings to set up a server as a syslog destination. For more information, see Configure syslog settings, on page 1297.

Step 9 Click **Syslog Servers** to specify the IP address, protocol used, format, and security zone for the syslog server that is designated as a logging destination. For more information, see Configure a syslog server, on page 1298.

Configure only a public syslog server. If you configure a local syslog server, health alerts sent to syslog are not delivered.

After you complete these steps, your Firewall Threat Defense device sends system events to the specified destinations for monitoring and analysis.

Firewall Threat Defense platform settings that apply to security event syslog messages

Security events include connection, Security Intelligence, intrusion, and file and malware events. Some syslog settings configured on the **Syslog** page and its tabs apply to syslog messages for security events. However, most settings only apply to messages for events related to system health and networking. This reference lists the settings that apply to syslog messages for security events.

Settings that apply to syslog messages for security events:

- **Logging Setup** tab:
  - **Send syslogs in EMBLEM format**
- **Syslog Settings** tab:
  - **Enable Timestamp on Syslog Messages**
  - **Timestamp Format**
  - **Enable Syslog Device ID**
- **Syslog Servers** tab:
  - All options on the **Add Syslog Server** form and the list of configured servers.

Configure logging and basic settings

Enable logging and configure syslog settings so the system generates and manages logs for data plane events. This setup ensures log data is archived and available for analysis or reporting.

Perform this task when you need to generate syslog messages for data plane events and manage log storage on your device. You can manipulate saved logging data. You can also specify actions for certain syslog messages, extract data for reporting, or track statistics using scripts.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense platform settings that apply to security event syslog messages , on page 1291.

Before you begin

- Determine where logs should be stored, such as an FTP server or flash memory.
- Understand the desired logging level (critical, alerts, emergencies, errors, etc)
- Identify if logging should be enabled on standby or failover units.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Logging Setup** in your device's configuration interface.

Step 3 Enable logging and configure basic logging settings.

- **Enable Logging**—Turns on the data plane system logging.
- **Enable Logging on the Failover Standby Unit**—Turns on logging for the standby device, if available.
- **Send syslogs in EMBLEM format**—Enables EMBLEM format logging for every logging destination. If you enable EMBLEM, you must use the UDP protocol to publish syslog messages. EMBLEM is not compatible with TCP.

  The EMBLEM syslog format is a Cisco-specific convention that is built upon the RFC 3164 and RFC 5424 standards. When EMBLEM is enabled, the syslog message prints a colon (:) after the <PRI> field, which differs from the RFC 5424 format.

  Note
  Syslog messages in RFC5424 format, typically displays the priority value (PRI). However, in Cloud-Delivered Firewall Management Center, if you want to display the PRI value in the syslog messages of the managed Firewall Threat Defense device, ensure to enable the EMBLEM format. Also, the device ID does not appear in EMBLEM-formatted syslog messages.

- **Send debug messages as syslogs**—Redirects debug trace output to the syslog. The syslog message number is 711001 and the default logging level is set to debug. To view these messages, ensure logging is enabled at the console and configured for the debug syslog message number and level.
- **Memory Size of Internal Buffer**—Specify the size of the internal buffer to which syslog messages are saved if the logging buffer is enabled. When the buffer fills up, it is overwritten. The default is 4096 bytes. The range is 4096 to 52428800 bytes.

Step 4 (Optional) Configure the syslog message logging to the Security Cloud Control.

a) Click the **All Logs** radio button to enable logging all the troubleshooting syslog messages corresponding to the selected severity level or click the **VPN Logs** radio button to enable logging only the VPN troubleshooting messages corresponding to the selected severity level.

b) Choose the syslog severity level for the logging messages in the **Logging Level** drop-down list.

- The logging level for **All Logs** is set to **critical** by default. You can choose to send syslog messages with severity levels **critical**, **alerts**, or **emergencies** to the Cloud-Delivered Firewall Management Center.
- The logging level for the VPN messages is set to **errors** by default.

  VPN troubleshooting syslogs can add excessive load on the Cloud-Delivered Firewall Management Center. Hence, enable this option with caution. Also, when you configure a device with site-to-site or remote access VPN, it automatically enables sending VPN syslogs to the management center by default. We recommend that you limit the logging level to **error** and above to restrict the excessive flow of syslogs to the Cloud-Delivered Firewall Management Center, especially in case of RAVPN, where multiple devices are involved.

For information on the levels, see Syslog message severity levels, on page 1284.

Step 5 (Optional) Save log buffer contents to an FTP server before the buffer is overwritten.

- Enable **FTP Server Buffer Wrap** and enter the necessary destination information in the following fields.
- **IP Address**—Select the host network object that contains the IP address of the FTP server.
- **User Name**—Enter the username to use when connecting to the FTP server.
- **Path**—Enter the path, relative to the FTP root, where the buffer contents should be saved.
- **Password/ Confirm**—Enter and confirm the password used to authenticate the username to the FTP server.

Step 6 (Optional) Save log buffer contents to flash memory before overwriting.

- Check the **Flash** checkbox
- Specify **Maximum flash to be used by logging (KB)** (range: 4-8044176 KB).
- Specify **Minimum free space to be preserved (KB)** (range: 0-8044176 kilobytes.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Logging is enabled and syslog messages are generated and stored based on your configuration. Log data is archived as specified and is accessible for further analysis or reporting.

Enable logging destinations

Enable logging destinations to ensure that syslog messages are received at the specified destination. This task allows you to control which messages are sent by applying message filters to each destination.

To see messages at a logging destination, enable the destination and specify a message filter.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense platform settings that apply to security event syslog messages , on page 1291.

Before you begin

Before you begin, ensure you have access to the system and the necessary permissions to configure logging destinations.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Logging Destinations**.

Step 3 Click **Add** to enable a destination and apply a logging filter, or edit an existing destination.

Step 4 In the **Logging Destinations** dialog box, select a destination and configure the filter for that destination:

a) Choose the destination to enable in the **Logging Destination** drop-down list. You can create one filter for each destination. Examples include Console, E-Mail, Internal buffer, SNMP trap, SSH Sessions, and Syslog servers.

Note
Console and SSH session logging works in the diagnostic CLI only. Enter **system support diagnostic-CLI**.

b) In **Event Class**, choose the filter that will apply to all classes not listed in the table.

You can configure these filters:

- **Filter on severity** —Select the severity level. Messages at this level or higher are sent to the destination
- **Use Event List** —Select the event list that defines the filter. You create these lists on the **Event Lists** page.
- **Disable Logging** —Prevents messages from being sent to this destination.

c) To create filters for each event class, click **Add** for a new filter or edit an existing filter. Select the event class and severity level to limit messages in that class. Click **OK** to save the filter.

For an explanation of the event classes, see Syslog message classes and associated message ID numbers, on page 1285.

d) Click **OK**.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

After completing this task, the selected logging destinations will receive syslog messages according to the filters you have configured.

Configure syslog messages to be sent to an e-mail address

Enable syslog messages to be delivered as e-mails to designated recipients so administrators receive important system notifications directly through e-mail.

You can set up a list of recipients for syslog messages to be sent as e-mails.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense platform settings that apply to security event syslog messages , on page 1291.

Before you begin

- Configure an SMTP server on the platform settings page
- Configure logging and basic settings, on page 1291
- Setup Email as the Logging Destination

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Email Setup**.

Step 3 Specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages.

Step 4 Click **Add** to enter a new e-mail address recipient of the specified syslog messages.

Step 5 Set the desired severity level for syslog messages for each recipient.

The severity filter ensures only messages of the specified level and higher are sent via e-mail. For information on the levels, see Syslog message severity levels, on page 1284.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Syslog messages are sent as e-mails to the specified recipients according to the configured severity level.

Create a custom event list

Create a custom event list to control which syslog messages are sent to a logging destination. This allows you to filter messages based on event class, severity, and message identifier for more granular control.

An event list is a custom filter you can apply to a logging destination, allowing greater control over which messages are sent. Normally, you filter messages for a destination based solely on severity. Using an event list, you can fine-tune filters using event class, severity, and message ID.

Creating a custom event list is a two-step process. You create a custom list in the **Event Lists**, and then use the event list to define the logging filter for the various types of destination, in the **Logging Destinations**.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense platform settings that apply to security event syslog messages , on page 1291.

Before you begin

Ensure you have access to the system where syslog event lists can be configured.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Events List**.

Step 3 Click **Add** to create a new list or click an existing list to edit.

a) Enter a unique name for the event list in the **Name** field; avoid spaces.

b) Set criteria to filter messages by severity or event class, select the **Severity/Event Class** tab and add or edit entries.

For information on the available classes see Syslog message classes and associated message ID numbers, on page 1285.

For information on the levels, see Syslog message severity levels, on page 1284.

Certain event classes do not apply to devices in transparent mode. If you configure these options, the system does not use or deploy them.

c) To identify messages specifically by message ID, select the **Message ID** and add or edit the IDs.

You can enter a range of IDs using a hyphen, for example, 100000-200000. Each message ID you enter must contain six digits. For information on how the initial three digits map to features, see Syslog message classes and associated message ID numbers, on page 1285.

For specific message numbers, see Cisco ASA Series Syslog Messages.

d) Click **OK** to save the event list.

Step 4 Click **Logging Destinations** and add or edit the destination that should use the custom event list filter.

See Enable logging destinations, on page 1293.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

The system creates the custom event list and applies it to the logging destination you select.

Configure syslog rate limits

Limit the rate at which syslog messages are generated, reducing log flooding and ensuring that important messages are not lost.

Syslog message generation can overwhelm logging systems if not managed properly. Configure rate limits for severity levels or specific message IDs to ensure efficient log management. When limits set for severity level and message ID conflict, the message ID limit takes precedence.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense platform settings that apply to security event syslog messages , on page 1291.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Rate Limit**.

Step 3 To limit message generation by severity level, click **Logging Level** > **Add**, then configure these options:

- **Logging Level**—Specify the severity level for the rate limit. For information on the levels, see Syslog message severity levels, on page 1284.
- **Number of messages**—The maximum number of messages of the specified type allowed in the specified time period.
- **Interval**—The number of seconds before the rate limit counter resets.

Step 4 Click **OK**.

Step 5 To limit message generation by syslog message ID, click **Syslog Level** > **Add**, then configure these options:

- **Syslog ID**—Specify the syslog message ID to be rate limited. For specific message numbers, see Cisco ASA Series Syslog Messages.
- **Number of messages**—Enter the maximum number of messages allowed in the specified interval.
- **Interval**—Set the interval (in seconds) for the rate limit counter to reset.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Syslog message generation is limited based on configured severity levels or message IDs, preventing excessive log messages and supporting efficient log management.

Configure syslog settings

Configure syslog settings to manage how syslog messages are generated and sent to external syslog servers. You can specify the facility code, timestamp, device ID, and control message severity and suppression.

General syslog settings let you choose what information is included in syslog messages sent to syslog servers. This information includes the facility code, timestamp, device ID, severity levels, and message suppression. If you configure devices to send syslog messages about security events (such as connection and intrusion events), some settings do not apply to messages that relate to security events. See *Threat Defense Platform Settings That Apply to Security Event Syslog Messages* in the *Cisco Secure Firewall Management Center Administration Guide*.

Before you begin

Make sure you have access to the device and the permissions needed to change syslog settings.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Syslog Settings**.

Step 3 Choose a system log facility from the **Facility** list as required by your syslog server.

Most UNIX systems use the default LOCAL4(20). Change the facility if your network devices use the same codes. In most cases, you do not need to set the facility value for security event logs.

Step 4 Select the **Enable timestamp on each syslog message** check box to include the date and time a message was generated in the syslog message.

Step 5 Select the desired **Timestamp Format**:

- **Legacy (MMM dd yyyy HH:mm:ss)**: Default format without a timezone (always UTC).
- **RFC 5424 (yyyy-MM-ddTHH:mm:ssZ)**: Uses ISO 8601 with “Z” indicating UTC.

Step 6 To add a device identifier at the start of messages, check the **Enable Syslog Device ID** check box and select the identifier type:

- **Interface**—Use the IP address of a specific interface.

  Note
  VTI tunnel interfaces are not supported for Syslog.

- **User Defined ID**—Enter a custom text string (up to 16 characters).
- **Host Name**—Use the device's hostname.

Step 7 Use the Syslog Message table to adjust severity or suppress messages:

Change a message's severity level or disable its generation if necessary. By default, syslog entries for NetFlow are enabled and appear in the table.

a) To suppress redundant messages due to NetFlow, select **Netflow Equivalent Syslogs**.

   This adds the messages to the table as suppressed messages.

   Note
   If any of these syslog equivalents are already in the table, your existing rules are not overwritten.

b) To add a rule, click **Add**.
c) Select the message number (Syslog ID) to change, from the **Syslog ID** drop-down list.
d) Choose the new severity level from the **Logging Level** drop-down list or select **Suppressed** to disable the generation of the message. Usually, you do not change both the severity level and disable the message. However, you can modify both fields if you want.
e) Click **OK** to add the rule to the table.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

- Deploy configuration changes.

Configure a syslog server

This task enables you to configure a syslog server for your device, allowing you to collect and monitor system-generated messages for security and operational visibility.

To configure a syslog server to handle messages generated from your system, perform this task.

If you want this syslog server to receive security events such as connection and intrusion events, see also Firewall Threat Defense platform settings that apply to security event syslog messages , on page 1291.

Note In version 7.4 and later, the Management and Diagnostic interfaces are merged. If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, you must use separate Platform Settings policies for merged and unmerged devices. This requirement also applies to version 7.3 and earlier, and to some Firewall Threat Defense devices upgraded to version 7.4.

Before you begin

- See requirements in Guidelines for logging, on page 1288.
- Make sure your devices can reach your syslog collector on the network.
- Ensure that only a public syslog server is configured. If you configure a local syslog server, health alerts will not be sent to syslog.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Syslog Server**.

Step 3 Check the **Allow user traffic to pass when TCP syslog server is down (Recommended)** check box, to allow traffic if any syslog server that is using the TCP protocol is down.

Note

- This option is enabled by default. Unless required, we recommend that you allow connections through the threat defense device when the external TCP syslog server is unreachable by the device.
- When the **Allow user traffic to pass when TCP syslog server is down** option is disabled in Cloud-Delivered Firewall Management Center version 6.2.x or earlier, it remains in the Disabled state even after upgrading to version 6.3 or later. You must manually enable this option.
- If this option is disabled and more than one TCP syslog server is configured on the device, user traffic is allowed if at least one server is reachable by thethreat defense device. This option only blocks traffic when none of the configured TCP syslog servers are reachable. The device generates a syslog message that describes the root cause of denied traffic.

```
%FTD-3-414003: TCP Syslog Server intf : IP_Address /port not responding. New connections
are denied based on logging permit-hostdown policy
```

Step 4 In the **Message queue size (messages)** field, enter the queue size for storing syslog messages on the security appliance when the syslog server is busy. The minimum is 1 message. The default is 512. Specify 0 to allow an unlimited number of messages to be queued (subject to available block memory).

If the number of messages exceeds the configured queue size, the excess messages are dropped, resulting in missing syslog entries. To determine the ideal queue size, you need to identify the available block memory. Use the **show blocks** command to know the current memory utilization. For more information on the command and its attributes, see *Cisco Secure Firewall ASA Series Command Reference Guide*. Contact Cisco TAC for assistance.

Step 5 Click **Add** to add a new syslog server.

a) In the **IP Address** drop-down list, select a network host object containing the IP address of the syslog server.

b) Choose the protocol, either TCP or UDP, and enter the port number used for communication between the Firewall Threat Defense device and the syslog server.

UDP is faster and uses fewer resources than TCP on the device.

The default port for UDP is 514. You must manually configure port 1470 for TCP. Valid non-default port values for either protocol are 1025 through 65535.

c) Check the **Log messages in Cisco EMBLEM format (UDP only)** check box to specify whether to log messages in Cisco EMBLEM format (available only if UDP is selected as the protocol).

The EMBLEM syslog format is a Cisco-specific convention that is built upon the RFC 3164 and RFC 5424 standards. Hence, when EMBLEM is enabled, the syslog message prints colon (:) after <PRI> field.

Note
Syslog messages in RFC5424 format typically displays the priority value, referred to as PRI. However, in Cloud-Delivered Firewall Management Center, only when you enable logging in Cisco EMBLEM format, the PRI value in the syslog messages of the managed Firewall Threat Defense is displayed. Also, the device ID does not appear in EMBLEM-formatted syslog messages.

d) Check the **Enable Secure Syslog** check box to encrypt the connection between the device and server using SSL or TLS over TCP.

Note
You must select TCP as the protocol and its port value ranging between 1025 and 65535 to use this option. Upload the certificate required to communicate with the syslog server on the **Devices** > **Certificates** page. Then, upload the certificate from the Firewall Threat Defense device to the syslog server to complete the secure relationship and allow it to decrypt the traffic. The **Enable Secure Syslog** option is not supported on the device Management interface.

e) Select **Device Management Interface** or **Security Zones or Named Interfaces** to communicate with the syslog server.

- **Device Management Interface**: Send syslogs out of the Management interface. We recommend that you use this option when configuring syslog on Snort events.

  Note
  The **Device Management Interface** option does not support the **Enable Secure Syslog** option.

- **Security Zones or Named Interfaces**: Select the interfaces from the list of **Available Zones** and click **Add**. You can also add virtual-router-aware interfaces.

  Important
  The Firewall Threat Defense data plane (Lina) syslog messages cannot be sent out through the diagnostic interface. Configure other interfaces or the Management interface (Br1/Management0) to send out the data plane syslog messages.

f) Click **OK**.

Step 6 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

- Deploy configuration changes.

Configure timeouts

Configure timeout values for console sessions and protocol connections in a Threat Defense policy to manage session durations and optimize device resource usage.

You can set the global idle timeout durations for the connection and translation slots of various protocols. If the slot has not been used for the idle time specified, the resource is returned to the free pool.

You can also set a timeout for console sessions on the device.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Timeouts**.

Step 3 Configure the timeouts you want to change.

For any setting, select **Custom** to specify your own value or **Default** to use the system default value. In most cases, the maximum timeout is 1193 hours.

You can disable some timeouts by selecting **Disable**.

- **Console Timeout**—The idle time until a connection to the console is closed, range is 0 or 5 to 1440 minutes. The default is 0, which means the session does not time out. If you change the value, existing console sessions will continue to use the previous timeout value. The updated timeout will apply only to new connections.
- **Translation Slot (xlate)**—The idle time until a NAT translation slot is freed. This duration must be at least 1 minute. The default is 3 hours.
- **Connection (Conn)**—The idle time until a connection slot is freed. This duration must be at least 5 minutes. The default is 1 hour.
- **Half-Closed**—The idle time until a TCP half-closed connection closes. A connection is considered half-closed if both the FIN and FIN-ACK have been seen. If only the FIN has been seen, the regular connection timeout applies. The minimum is 30 seconds. The default is 10 minutes.
- **UDP**—The idle time until a UDP connection closes. This duration must be at least 1 minute. The default is 2 minutes.
- **ICMP**—The idle time after which general ICMP states are closed. The default (and minimum) is 2 seconds.

- **RPC/Sun RPC**—The idle time until a SunRPC slot is freed. This duration must be at least 1 minute. The default is 10 minutes.

  In a Sun RPC-based connection, if the parent connection is deleted or times out, a new child connection might not be treated as part of the parent–child relationship. As a result, the new connection could be evaluated according to the system's policy or rules. After the parent connection times out, existing child connections remain valid only until their own timeout value is reached.

- **H.225**—The idle time until an H.225 signaling connection closes. The default is 1 hour. To close a connection immediately after all calls are cleared, a timeout of 1 second (0:0:1) is recommended.
- **H.323**—The idle time after which H.245 (TCP) and H.323 (UDP) media connections close. The default (and minimum) is 5 minutes. Because the same connection flag is set on both H.245 and H.323 media connections, the H.245 (TCP) connection shares the idle timeout with the H.323 (RTP and RTCP) media connection.
- **SIP**—The idle time until a SIP signaling port connection closes. This duration must be at least 5 minutes. The default is 30 minutes.
- **SIP Media**—The idle time until a SIP media port connection closes. This duration must be at least 1 minute. The default is 2 minutes. The SIP media timer is used for SIP RTP/RTCP with SIP UDP media packets, instead of the UDP inactivity timeout.
- **SIP Disconnect**—The idle time after which SIP session is deleted if the 200 OK is not received for a CANCEL or a BYE message, between 0:0:1 and 0:10:0. The default is 2 minutes (0:2:0).
- **SIP Invite**—The idle time after which pinholes for PROVISIONAL responses and media xlates will be closed, between 0:1:0 and 00:30:0. The default is 3 minutes (0:3:0).
- **SIP Provisional Media**—The timeout value for SIP provisional media connections, between 1 and 30 minutes. The default is 2 minutes.
- **Floating Connection**—When multiple routes exist to a network with different metrics, the system uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To make it possible to use better routes, set the timeout to a value between 0:0:30 and 1193:0:0. This timer does not apply to connections through virtual tunnel interfaces (VTI). If a connection through a VTI gets stuck, you must manually clear it.
- **Xlate PAT**—The idle time until a PAT translation slot is freed, between 0:0:30 and 0:5:0. The default is 30 seconds. You may want to increase the timeout if upstream routers reject new connections using a freed PAT port because the previous connection might still be open on the upstream device.
- **TCP Proxy Reassembly**—The idle timeout after which buffered packets waiting for reassembly are dropped, between 0:0:10 and 1193:0:0. The default is 1 minute (0:1:0).
- **ARP Timeout**—The number of seconds between ARP table rebuilds, from 60 to 4294967. The default is 14,400 seconds (4 hours).

Step 4 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

The configured timeout values are applied to the Threat Defense policy. Sessions and protocol connections will now use the specified timeout durations after you deploy the policy.

Synchronize time on FTD devices

Use a Network Time Protocol (NTP) server to synchronize the clock settings on your devices. We recommend you configure all Firewall Threat Defenses managed by a Cloud-Delivered Firewall Management Center to use the same NTP server as the Cloud-Delivered Firewall Management Center. The Firewall Threat Defense gets its time directly FROM the configured NTP server. If the Firewall Threat Defense's configured NTP servers are not reachable for any reason, it synchronizes its time with the Cloud-Delivered Firewall Management Center.

The device supports NTPv4.

Note If you are deploying Firewall Threat Defense on the Firepower 4100/9300 chassis, you must configure NTP on the Firepower 4100/9300 chassis so that Smart Licensing will work properly and to ensure proper timestamps on device registrations. You should use the same NTP server for the Firepower 4100/9300 chassis and the Cloud-Delivered Firewall Management Center.

Before you begin

- If your organization has one or more NTP servers that your Firewall Threat Defense can reach, use the same NTP server or servers for your devices that you have configured for Time Synchronization on the **System** (⚙) > **Configuration** page on your Cloud-Delivered Firewall Management Center.
- If you selected **Use the authenticated NTP server only** when configuring NTP server or servers for the Cloud-Delivered Firewall Management Center, for your devices use only the NTP server or servers that are configured to authenticate with the Cloud-Delivered Firewall Management Center. (The managed devices will use the same NTP servers as the Cloud-Delivered Firewall Management Center, but their NTP connections will not use authentication.)
- If your device cannot reach an NTP server or your organization does not have one, use the **Via NTP FROM Management Center** option as discussed in this procedure.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Time Synchronization**.

Step 3 Configure one of these clock options:

- **Via NTP FROM Management Center**—(Default). The managed device gets time FROM the NTP servers you configured for the Cloud-Delivered Firewall Management Center (except for authenticated NTP servers) and synchronizes time with those servers directly. However, if any of these are true, the managed device synchronizes time FROM the Cloud-Delivered Firewall Management Center:
  - The Cloud-Delivered Firewall Management Center's NTP servers are not reachable by the device.
  - The Cloud-Delivered Firewall Management Center has no unauthenticated servers.

- **Via NTP FROM**—If your Cloud-Delivered Firewall Management Center is using NTP servers on the network, select this option and enter the fully-qualified DNS name (such as NTP.example.com), or IPv4 or IPv6 address, of the same NTP servers you specified in **System** (⚙) > **Configuration** > **Time Synchronization**. If the NTP servers are not reachable, the Cloud-Delivered Firewall Management Center acts as an NTP server.

When multiple NTP servers are configured, the device uses the NTP server that is deemed appropriate based on the criteria defined in RFC. Thus, the status "Being used" for a specific NTP server indicates that the server is currently used by the device.

Step 4 Click **Save**.

After completing these steps, the FTD device synchronizes its clock with the configured NTP server or, if unavailable, with the management center, ensuring accurate time for logs and system operations.

What to do next

- Deploy configuration changes.

Configure a time zone for a device

By default, your device uses the UTC time zone. To set a different time zone for the device, use this procedure.

The time zone you specify is used only in policies that use time-based policy application.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit a Firewall Threat Defense policy.

You can also create time zone objects from the **Objects** > **Object Management** > **Time Zone** page.

Step 2 Create a new time zone object by clicking **+**.

Step 3 Select the time zone.

Step 4 Click **Save**.

What to do next

- Create time range objects, select applicable time ranges in access control and prefilter rules, and assign the parent policies to devices associated with the correct time zone.
- Deploy configuration changes.

Configure UCAPL/CC compliance

Configure Unified Capabilities Approved Products List (UCAPL) or Common Criteria (CC) compliance to ensure your Secure Firewall Threat Defense device meets security certifications and export-controlled feature requirements.

Use this task when you need to enable security certifications compliance for your device. This is required for export-controlled features and may be necessary for regulatory or organizational compliance.

For more information about this setting and how to enable it for the Cloud-Delivered Firewall Management Center, see the *Cisco Secure Firewall Management Center Administration Guide*.

Caution After you enable this setting, disabling it is not possible. To return the appliance to standard mode, reimage the appliance.

Before you begin

- Secure Firewall Threat Defense devices cannot use an evaluation license; your Smart Software Manager account must be enabled for export-controlled features.
- Secure Firewall Threat Defense devices must be deployed in routed mode.
- You must have administrative privileges to perform this task.
- Security certifications compliance mode cannot be changed if the Firewall Threat Defense device is in high availability. Change the security certifications compliance mode before you create the high availability pair.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **UCAPL/CC Compliance**.

Step 3 To *permanently* enable security certifications compliance on the appliance, you have two choices:

- To enable security certifications compliance in Common Criteria mode, choose **CC** from the drop-down list.
- To enable security certifications compliance in Unified Capabilities Approved Products List (UCAPL) mode, choose **UCAPL** from the drop-down list.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

After completing these steps, your device operates in the selected compliance mode (UCAPL or CC), and the policy is enforced once deployed.

Configure a performance profile

Configure the performance profile to control how CPU cores are allocated between the data plane and Snort processes on supported devices. This configuration optimizes device performance for VPN, improves intrusion prevention, or balances the workload.

The performance profile determines how the CPU cores on the device are assigned to two of the main system processes: the data plane (Lina) and Snort. The data plane handles VPN connections, routing, and basic layer 3/4 processing. Snort provides advanced inspection for intrusion and malware prevention, URL filtering, application filtering, and other deep packet inspection features.

If you use both basic and advanced features, keep the performance profile unchanged. The system is designed to provide a balanced assignment of cores to these processes. This assignment differs based on the hardware model.

However, if you use the device primarily for VPN, or for intrusion and other advanced inspection, you can skew the performance profile so that more cores are assigned to the more heavily used features. This change can improve system performance.

Before you begin

- These settings apply to systems running release 7.3+ only.
- Performance profile is supported on these device types:
  - Firepower 4100/9300
  - Secure Firewall 3100/4200 (7.4+)
  - Secure Firewall Threat Defense Virtual
- Changing the performance profile is not supported on units in a cluster or high-availability group, or those configured for multi-instance.
- The minimum number for core allocation is 2. Cores are assigned in even numbers based on the selected performance profile.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Performance Profile**.

Step 3 Select a profile:

- **Default**—This is the recommended setting and is the best option if you configure both VPN and intrusion inspection.
- **VPN Heavy with prefilter fastpath**—If you primarily use the device as a VPN endpoint or headend, and you configure rules in the prefilter policy to fastpath VPN traffic, you can choose this option to assign the majority of CPU cores to the data plane. The allocation is 90% data plane, 10% Snort.
- **VPN Heavy with inspection**—If you primarily use the device as a VPN endpoint or headend, but do not use the prefilter policy to fastpath VPN traffic, you can choose this option to assign the majority of

CPU cores to the data plane. This option assumes that you leave intrusion inspection, URL filtering, and other advanced functions that use Snort, to a different device in the network. The allocation is 60% data plane, 40% Snort.

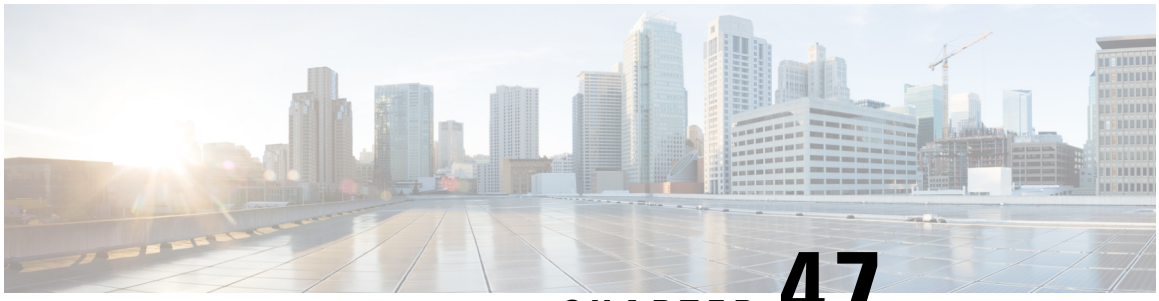
- **IPS Heavy**—If you do not configure VPN, but you do use the device for intrusion prevention, you can choose this option to assign the majority of CPU core to the Snort process. The allocation is 30% data plane, 70% Snort.

Step 4 Click **Save**.

Step 5 Deploy the policy to apply your configuration.

Step 6 After deployment completes, reboot each affected device to apply the new core assignments.

After completing these steps, the device will use the selected performance profile and allocate CPU cores according to your configuration. The new core assignments take effect after you reboot the device.



CHAPTER 47

FlexConfig policies

These topics describe how to configure and deploy FlexConfig policies.

- FlexConfig policy overview, on page 1309
- Requirements and prerequisites for FlexConfig policies, on page 1330
- Guidelines and limitations for FlexConfig, on page 1331
- Customize device configuration with FlexConfig policies, on page 1331
- Examples for FlexConfig, on page 1346
- Migrate FlexConfig policies, on page 1352

FlexConfig policy overview

A FlexConfig policy is a container of an ordered list of FlexConfig objects that

- includes a series of Apache Velocity scripting language commands, ASA software configuration commands, and variables that you define
- contains FlexConfig objects that are essentially programs that generate a sequence of ASA commands that will then be deployed to the assigned devices, and
- configures the related feature on the Firewall Threat Defense device.

FlexConfig policy usage and limitations

Firewall Threat Defense uses ASA configuration commands to implement some features, but not all features. There is no unique set of Firewall Threat Defense configuration commands. Instead, the point of FlexConfig is to allow you to configure features that are not yet directly supported through Cloud-Delivered Firewall Management Center policies and settings.

Caution We recommend using FlexConfig policies only if you are an advanced user with a strong ASA background and at your own risk. You may configure any commands that are not prohibited. Enabling features through FlexConfig policies may cause unintended results with other configured features.

Contact the Cisco Technical Assistance Center for support concerning FlexConfig policies that you have configured. The Cisco Technical Assistance Center does not design or write custom configurations on any customer's behalf. Cisco expresses no guarantees for correct operation or interoperability with other features. FlexConfig features may become deprecated at any time. For fully guaranteed feature support, you must wait for the Cloud-Delivered Firewall Management Center support. When in doubt, do not use FlexConfig policies.

Recommended usage for FlexConfig policies

Use FlexConfig policies in two main scenarios to ensure proper firewall configuration and functionality.

- You are converting from ASA to Firewall Threat Defense, and there are compatible features you are using (and need to continue using) that Cloud-Delivered Firewall Management Center does not directly support. In this case, use the **show running-config** command on the ASA to see the configuration for the feature and create your FlexConfig objects to implement it. Experiment with the object's deployment settings (once/every time and append/prepend) to get the right setting. Verify by comparing **show running-config** output on the two devices.
- You are using Firewall Threat Defense but there is a setting or feature that you need to configure, e.g. the Cisco Technical Assistance Center tells you that a particular setting should resolve a specific problem you are encountering. For complicated features, use a lab device to test the FlexConfig and verify that you are getting the expected behavior.

The system includes a set of predefined FlexConfig objects that represent tested configurations. If the feature you need is not represented by these objects, first determine if you can configure an equivalent feature in standard policies. For example, the access control policy includes intrusion detection and prevention, HTTP and other types of protocol inspection, URL filtering, application filtering, and access control. The ASA implements these functions using separate features. Many features are not configured using CLI commands, so not every policy appears in the output of the **show running-config** command.

Note At all times, keep in mind that there is not a one-to-one overlap between ASA and Firewall Threat Defense. Do not attempt to completely recreate an ASA configuration on a Firewall Threat Defense device. You must carefully test any feature that you configure using FlexConfig.

CLI commands in FlexConfig objects

A managed device uses ASA configuration commands to configure some features. Although not all ASA features are compatible with the device, some features can work on the managed device but you cannot configure them through Cloud-Delivered Firewall Management Center policies. You can use FlexConfig objects to specify the CLI required to configure these features.

ASA documentation resources

When using FlexConfig to manually configure a feature, you are responsible for knowing and implementing the commands according to the proper syntax. FlexConfig policies do not validate CLI command syntax. For more information about proper syntax and configuring CLI commands, use the ASA documentation as a reference:

- ASA CLI configuration guides explain how to configure a feature. Find the guides at https://www.cisco.com/c/en/us/support/security/adaptive-security-appliance-asa-software/products-installation-and-configuration-guides-list.html
- ASA command references provide additional information sorted by command name. Find the references at https://www.cisco.com/c/en/us/support/security/adaptive-security-appliance-asa-software/products-command-reference-list.html

Determine the ASA software version and current CLI configuration

Determine the ASA software version and examine the current CLI-based configuration to identify which ASA CLI configuration guides to use for instructions on configuring a feature.

Because the system uses ASA software commands to configure some features, you need to determine the current ASA version used in software running on the Firewall Threat Defense device. This version number indicates which ASA CLI configuration guides to use for instructions on configuring a feature. You also should examine the current CLI-based configuration and compare it to the ASA configuration you want to implement.

Keep in mind that any ASA configuration will be very different from a Firewall Threat Defense configuration. Many Firewall Threat Defense policies are configured outside of the CLI, so you cannot see the configuration by looking at the commands. Do not try to create a one-to-one correspondence between an ASA and Firewall Threat Defense configuration.

To view this information, make an SSH connection to the device's management interface and issue the following commands:

- **show version system** and look for the Cisco Adaptive Security Appliance Software Version number. (If you issue the command through the Cloud-Delivered Firewall Management Center CLI tool, omit the **system** keyword.)
- **show running-config** to view the current CLI configuration.
- **show running-config all** to include all the default commands in the current CLI configuration.

You can also issue these commands from within Cloud-Delivered Firewall Management Center using this procedure.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 Click the name of the device targeted by the FlexConfig policy.

You might need to click the open/close arrow in the **Count** column in the Status table to see any devices.

Step 3 Click **View System and Troubleshoot Details**.

Step 4 Click **Advanced Troubleshooting**.

Step 5 Click **Threat Defense CLI**.

Step 6 Choose the **Device**, then choose **show** as the command, and type **version** or one of the other commands as the parameter.

Step 7 Click **Execute**.

For version, search the output for the Cisco Adaptive Security Appliance Software Version number.

You can select the output and press Ctrl+C, then paste it into a text file for later analysis.

Prohibited CLI commands

The purpose of FlexConfig is to configure features that are available on ASA devices that you cannot configure on Firewall Threat Defense devices using Cloud-Delivered Firewall Management Center.

Thus, you are prevented from configuring ASA features that have equivalents in Cloud-Delivered Firewall Management Center. The following table lists some of these prohibited command areas.

In addition, some **clear** commands are prohibited because they overlap with managed policies, and can delete part of the configuration for a managed policy.

The FlexConfig object editor prevents you from including prohibited commands in the object.

Prohibited CLI Command	Description
AAA	Configuration blocked.
AAA-Server	Configuration blocked.
Access-list	Advanced ACL, Extended ACL, and Standard ACL are blocked. Ethertype ACL is allowed. You can use standard and extended ACL objects defined in the object manager inside the template as variables.
ARP Inspection	Configuration blocked.
As-path Object	Configuration blocked.
Banner	Configuration blocked.
BGP	Configuration blocked.
Clock	Configuration blocked.
Community-list Object	Configuration blocked.
Copy	Configuration blocked.
Delete	Configuration blocked.
DHCP	Configuration blocked.
Enable Password	Configuration blocked.

Prohibited CLI Command	Description
Erase	Configuration blocked.
Fragment Setting	Blocked, except for **fragment reassembly**.
Fsck	Configuration blocked.
HTTP	Configuration blocked.
ICMP	Configuration blocked.
Interface	Only **nameif**, **mode**, **shutdown**, **ip address** and **mac-address** commands are blocked.
Multicast Routing	Configuration blocked.
NAT	Configuration blocked.
Network Object/Object-group	Network object creation in the FlexConfig object is blocked, but you can use network objects and groups defined in the object manager inside the template as variables.
NTP	Configuration blocked.
OSPF/OSPFv3	Configuration blocked.
pager	Configuration blocked.
Password Encryption	Configuration blocked.
Policy-list Object	Configuration blocked.
Prefix-list Object	Configuration blocked.
Reload	You cannot schedule reloads. The system does not use the **reload** command to restart the system, it uses the **reboot** command.
RIP	Configuration blocked.
Route-Map Object	Route-map object creation in the FlexConfig object is blocked, but you can use route map objects defined in the object manager inside the template as variables.
Service Object/Object-group	Service object creation in the FlexConfig object is blocked, but you can use port objects defined in the object manager inside the template as variables.
SNMP	Configuration blocked.
SSH	Configuration blocked.
Static Route	Configuration blocked.
Syslog	Configuration blocked.

Prohibited CLI Command	Description
Time Synchronization	Configuration blocked.
Timeout	Configuration blocked.
VPN	Configuration blocked.

Template scripts

You can use scripting language to control processing within a FlexConfig object. Scripting language instructions are a subset of commands supported in the Apache Velocity 1.3.1 template engine. The Apache Velocity engine is a Java-based scripting language that supports looping, if/else statements, and variables.

To learn how to use the scripting language, see the *Velocity Developer Guide* at http://velocity.apache.org/engine/devel/developer-guide.html.

FlexConfig variables

You can use variables in a FlexConfig object in cases where part of a command or processing instruction depends on runtime information rather than static information. During deployment, the variables are replaced with strings obtained from other device configurations, depending on the type of variable.

FlexConfig variable types

Variables are replaced with strings based on the type of variable.

- Policy object variables are replaced with strings obtained from objects defined in Cloud-Delivered Firewall Management Center.
- System variables are replaced with information obtained from the device itself or from policies configured for it.
- Processing variables are loaded with the contents of policy object or system variables as during script execution. For example, in a loop, you iteratively load one value from a policy object or system variable into a processing variable, then use the processing variable to form a command string or perform another action. These processing variables do not show up in the Variables list within a FlexConfig object. Also, you do not add them using the **Insert** menu in the FlexConfig object editor.
- Secret key variables are replaced with the single string defined for the variable within the FlexConfig object.

Variables start with the $ character, except for secret keys, which start with the @ character. For example, $ifname is a policy object variable in this command, whereas @keyname is a secret key.

```
interface $ifname
key @keyname
```

Note The first time you insert a policy object or system variable, you must do so through the **Insert** menu in the FlexConfig object editor. This action adds the variable to the **Variables** list at the bottom of the FlexConfig object editor. You must type in the variable string on subsequent uses, even when using system variables. If you are adding a processing variable, which does not have an object or system variable assignment, do not use the **Insert** menu. If you are adding a secret key, always use the **Insert** menu. Secret key variables do not show up in the Variables list.

Whether a variable is resolved as a single string, a list of strings, or a table of values depends on the type of policy object or system variable you assign to the variable. (Secret keys always resolve to a single string.) You must understand what will be returned in order to process the variables correctly.

Variable processing

At runtime, a variable can resolve to any of these: a single string; a list of strings of the same type; a list of strings of different types; or a table of named values. Variables that resolve to multiple values can have a determinate or indeterminate length. To process the values correctly, you must understand what will be returned.

The main possibilities include these scenarios.

Single value variables

If a variable always resolves to a single string, use the variable directly without modification in the FlexConfig script.

For example, the predefined text variable tcpMssBytes always resolves to a single value (which must be numeric). The **Sysopt_basic** FlexConfig then uses an if/then/else structure to set the maximum segment size based on the value of another single-value text variable, *tcpMssMinimum*.

Implementation methods

The implementation of single value variables depends on the variable type:

- **Standard text variables**: Use the variable directly in the FlexConfig script. For example, the predefined text variable tcpMssBytes always resolves to a single numeric value.
- **Secret key variables**: Always use the **Insert** menu to add the variable, even for second and subsequent uses. These variables do not show up in the Variables list within the FlexConfig object.
- **Policy object variables for network objects**: Equate to a single IP address specification (host address, network address, or address range). You must know what type of address to expect because ASA commands require specific address types.

Note Policy object variables for network objects also equate to a single IP address specification, either a host address, network address, or address range. However, in this case, you must know what type of address to expect, because the ASA commands require specific address types. For example, if a command requires a host address, using a network object variable that points to an object that contains a network address will result in an error during deployment.

TCP MSS configuration using single value variables

The **Sysopt_basic** FlexConfig uses an if/then/else structure to set the maximum segment size based on the value of another single-value text variable, tcpMssMinimum:

```
#if($tcpMssMinimum == "true")
  sysopt connection tcpmss minimum $tcpMssBytes
#else
  sysopt connection tcpmss $tcpMssBytes
#end
```

In this example, you would use the **Insert** menu in the FlexConfig object editor to add the first use of $tcpMssBytes, but you would type in the variable directly on the #else line.

Multiple value variables with same type values

Several policy object and system variables resolve to multiple values of the same type. For example, an object variable that points to a network object group resolves to a list of the IP addresses within the group. Similarly, the system variable $SYS_FW_INTERFACE_NAME_LIST resolves to a list of interface names.

You can also create text objects for multiple values of the same type. For example, the predefined text object enableInspectProtocolList can contain more than one protocol name.

Processing characteristics

Multiple value variables that resolve to a list of items of the same type are frequently of indeterminate length. For example, you cannot know beforehand how many interfaces on a device are named, as users can configure or unconfigure interfaces at any time.

You would typically use a loop to process multiple value variables of the same type. For example, the predefined FlexConfig **Default_Inspection_Protocol_Enable** uses a #foreach loop to go through the enableInspectProtocolList object and process each value.

```
policy-map global_policy
    class inspection_default
    #foreach ( $protocol in $enableInspectProtocolList)
    inspect $protocol
    #end
```

In this example, the script assigns each value in turn to the $protocol variable, which is then used in an ASA **inspect** command to enable the inspection engine for that protocol. In this case, you simply type in $protocol as a variable NAME. You do not use the **Insert** menu to add it, because you are not assigning an object or system value to the variable. However, you must use the **Insert** menu to add $enableInspectProtocolList.

The system loops through the code between #foreach and #end until there are no values remaining in $enableInspectProtocolList.

Multiple value variables with different types

You can create multiple value text objects and have each value serve a different purpose. For example, the predefined **netflow_Destination** text object should include three values in this order: interface name, destination IP address, and UDP port number.

Objects defined in this way should have a determinate number of values to make them easier to process.

Processing multiple value variables

Use the get method to process these objects. Type **.get(*n*)** at the end of the object name, replacing *n* with an index into the object. Start counting at 0, even though the text object lists its values starting at 1.

NetFlow destination configuration example

The predefined **netflow_Destination** text object should have 3 values, in order, interface name, destination IP address, and UDP port number.

The Netflow_Add_Destination object uses this line to add the 3 values from netflow_Destination to the ASA **flow-export** command:

```
flow-export destination $netflow_Destination.get(0) $netflow_Destination.get(1)
$netflow_Destination.get(2)
```

In this example, you would use the **Insert** menu in the FlexConfig object editor to add the first use of $netflow_Destination, and then add **.get(0)**. But you would type in the variable directly for the **$netflow_Destination.get(1)** and **$netflow_Destination.get(2)** specifications.

Multiple value variables that resolve to a table of values

Some system variables return a table of values. These variables include the string MAP in their names. For example, $SYS_FTD_ROUTED_INTF_MAP_LIST is one such variable. The routed interface map returns data in a table of values. Line breaks have been added for clarity.

Variable structure and processing

These variables return data that includes a table of named values for each item. For example, the ROUTED interface MAP returns data that looks like this (line returns added for clarity):

```
[{intf_hardwarare_id=GigabitEthernet0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.10.1, intf_ipv6_link_local_address=,
intf_logical_name=outside},

{intf_hardwarare_id=GigabitEthernet0/1, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.11.1, intf_ipv6_link_local_address=,
intf_logical_name=inside},

{intf_hardwarare_id=GigabitEthernet0/2, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=},

{intf_hardwarare_id=Management0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=management}]
```

In the above example, information is returned for 4 interfaces. Each interface includes a table of named values. For example, intf_hardwarare_id is the name of the interface hardware name property, and returns strings such as GigabitEthernet0/0.

This type of variable is typically of indeterminate length, so you need to use looping to process the values. You also need to add the property name to the variable name to indicate which value to retrieve.

IS-IS interface configuration

For example, IS-IS configuration requires that you add the ASA **ISIS** command to an interface that has a logical name in interface configuration mode. However, you enter that mode using the interface's hardware name. Thus, you need to identify which interfaces have logical names, then configure just those interfaces using their hardware names. The ISIS_Interface_Configuration predefined FlexConfig does this using an if/then structure nested in a loop. In this code, you can see that the #foreach scripting command loads each interface MAP into the $INTF variable, then the #if statement keys off the intf_logical_name value in the MAP ($INTF.intf_logical_name), and if the value is in the LIST defined in the isisIntfList predefined text variable, enters the interface command using the intf_hardwarare_id value ($INTF.intf_hardwarare_id). You would need to edit the isisIntfList variable to add the names of the interfaces on which to configure IS-IS.

```
#foreach ($intf in $SYS_FTD_ROUTED_INTF_MAP_LIST)
 #if ($isIsIntfList.contains($intf.intf_logical_name))
  interface $intf.intf_hardwarare_id
   isis
   #if ($isIsAddressFamily.contains("ipv6"))
   ipv6 router isis
   #end
 #end
#end
```

Evaluate device variable results

To evaluate what a variable will return by creating a simple FlexConfig object that processes variables and shows their resolved values in a configuration preview.

This method allows you to test and analyze variable resolution without deploying potentially invalid configuration commands to devices.

Procedure

Step 1 Create a simple FlexConfig object that processes an annotated list of variables.

Example:

```
Following is a network object group variable for the
IPv4-Private-All-RFC1918 object:

$IPv4_Private_addresses

Following is the system variable SYS_FW_MANAGEMENT_IP:

$SYS_FW_MANAGEMENT_IP

Following is the system variable SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST:

$SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST

Following is the system variable SYS_FTD_ROUTED_INTF_MAP_LIST:

$SYS_FTD_ROUTED_INTF_MAP_LIST
```

```
Following is the system variable SYS_FW_INTERFACE_NAME_LIST:

$SYS_FW_INTERFACE_NAME_LIST
```

Step 2 Assign the FlexConfig object to a FlexConfig policy.

Step 3 Assign the policy to a device.

Step 4 Save the policy.

Step 5 Preview the configuration for that device.

The preview shows the resolved values.

Step 6 Select the preview text, copy it, then paste the output into a text file for analysis.

Example:

```
###Flex-config Prepended CLI ###

###CLI generated from managed features ###

###Flex-config Appended CLI ###
Following is an network object group variable for the
IPv4-Private-All-RFC1918 object:

[10.0.0.0, 172.16.0.0, 192.168.0.0]

Following is the system variable SYS_FW_MANAGEMENT_IP:

192.168.0.171

Following is the system variable SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST:

[dns, ftp, h323 h225, h323 ras, rsh, rtsp, sqlnet, skinny, sunrpc,
xdmcp, sip, netbios, tftp, icmp, icmp error, ip-options]

Following is the system variable SYS_FTD_ROUTED_INTF_MAP_LIST:

[{intf_hardwarare_id=GigabitEthernet0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.10.1, intf_ipv6_link_local_address=,
intf_logical_name=outside},

{intf_hardwarare_id=GigabitEthernet0/1, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.11.1, intf_ipv6_link_local_address=,
intf_logical_name=inside},

{intf_hardwarare_id=GigabitEthernet0/2, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=},

{intf_hardwarare_id=Management0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=management}]

Following is the system variable SYS_FW_INTERFACE_NAME_LIST:

[outside, inside, management]
```

Note

Do not deploy this FlexConfig to the device, however, because it will not contain any valid configuration commands. You would get deployment errors. After obtaining the preview, delete the FlexConfig object from the FlexConfig policy and save the policy.

You can analyze the resolved variable values in the text file to understand what each variable returns for the device.

FlexConfig policy object variables

A policy object variable is associated with a specific policy object configured in the Object Manager. When you insert a policy object variable in a FlexConfig object, you give the variable a name and select the object associated with it.

Policy object variable types

Although you can give the variable the exact same name as the associated object, the variable itself is not the same thing as the associated object. You must use the **Insert** > **Insert Policy Object** > ***Object Type*** menu in the FlexConfig object editor to add the variable for the first time to the script in the FlexConfig to establish the association with the object. Simply typing in the name of the object preceded by a $ sign does not create a policy object variable.

You can create variables to point to these types of objects. Ensure that you create the right type of object for each variable. To create objects, go to the **Objects** > **Object Management** page.

- **Text Objects**—For text strings, which can include IP addresses, numbers, and other free-form text such as interface or zone names. Select **FlexConfig** > **Text Object** from the table of contents, then click **Add Text Object**. You can configure these objects to contain a single value or multiple values. These objects are highly flexible and built specifically for use within FlexConfig objects. For detailed information, see Configure FlexConfig text objects, on page 1337.
- **Network**—For IP addresses. You can use network objects or groups. Select **Network** from the table of contents, then select **Add Network** > **Add Object** or **Add Group**. If you use a group object, the variable returns a list of each IP address specification within the group. Addresses can be host, network, or address ranges, depending on the object contents. See Networks, on page 2982.
- **Security Zones**—For interfaces within a security zone or interface group. Select **Interface** from the table of contents, then select **Add** > **Security Zone** or **Interface Group**. A security zone variable returns a list of the interfaces within that zone or group for the device being configured. See Interfaces, on page 2980.
- **Standard ACL Object**—For standard access control lists. A standard ACL variable returns the name of the standard ACL object. Select **Access List** > **Standard** from the table of contents, then click **Add Standard Access List Object**. See Access lists, on page 2948.
- **Extended ACL Object**—For extended access control lists. An extended ACL variable returns the name of the extended ACL object. Select **Access List** > **Extended** from the table of contents, then click **Add Extended Access List Object**. See Access lists, on page 2948.
- **Route Map**—For route map objects. A route map variable returns the name of the route map object. Select **Route Map** from the table of contents, then click **Add Route Map**. See Configure a route map, on page 3017.

FlexConfig system variables

System variables are replaced with information obtained from the device itself or from policies configured for it. To add a variable for the first time in the FlexConfig script and establish its association with the system variable, use the **Insert** > **Insert System Variable** > ***Variable NAME*** menu in the FlexConfig object editor. Typing the system variable name with a preceding $ sign does not create a system variable in the context of the FlexConfig object.

Before using a variable, check the typical value it returns. Refer to Evaluate device variable results, on page 1318.

Table 89: System variables

NAME	Description
SYS_FW_OS_MODE	The operating system MODE of the device. Possible values are ROUTED or TRANSPARENT.
SYS_FW_OS_MULTIPLICITY	Whether the device is running in SINGLE or multiple context MODE. Possible values are SINGLE, MULTI, or NOT_APPLICABLE.
SYS_FW_MANAGEMENT_IP	The MANAGEMENT IP address of the device
SYS_FW_HOST_NAME	The device hostname
SYS_FTD_INTF_POLICY_MAP	A MAP with INTERFACE NAME as key and POLICY-MAP as value. This variable returns nothing if there are no INTERFACE-based service policies defined on the device.
SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST	The LIST of protocols for which inspection is ENABLED.
SYS_FTD_ROUTED_INTF_MAP_LIST	A LIST of ROUTED INTERFACE maps on the device. Each MAP includes a set of named values related to ROUTED INTERFACE configuration.
SYS_FTD_SWITCHED_INTF_MAP_LIST	A LIST of SWITCHED INTERFACE maps on the device. Each MAP includes a set of named values related to SWITCHED INTERFACE configuration.
SYS_FTD_INLINE_INTF_MAP_LIST	A LIST of INLINE INTERFACE maps on the device. Each MAP includes a set of named values related to INLINE set INTERFACE configuration.
SYS_FTD_PASSIVE_INTF_MAP_LIST	A LIST of PASSIVE INTERFACE maps on the device. Each MAP includes a set of named values related to PASSIVE INTERFACE configuration.
SYS_FTD_INTF_BVI_MAP_LIST	A LIST of Bridge Virtual INTERFACE maps on the device. Each MAP includes a set of named values related to BVI configuration.
SYS_FW_INTERFACE_HARDWARE_ID_LIST	A LIST of the HARDWARE names for interfaces on the device, such as GigabitEthernet0/0.

NAME	Description
SYS_FW_INTERFACE_NAME_LIST	A LIST of logical names for interfaces on the device, such as inside.
SYS_FW_INLINE_INTERFACE_NAME_LIST	A LIST of logical names for interfaces configured as PASSIVE or ERSPAN PASSIVE.
SYS_FW_NON_INLINE_INTERFACE_NAME_LIST	A LIST of logical names for interfaces that are NOT part of INLINE sets, such as all ROUTED interfaces.

Predefined FlexConfig objects

The predefined FlexConfig objects provide tested configurations for select features. Use these objects if you need to configure these features, which otherwise cannot be configured using the management center. Note that some objects might become obsolete as features are added to the GUI.

This table lists the available objects. Make note of the associated text objects. You must edit these text objects to customize the behavior of the predefined FlexConfig object. The text objects make it possible for you to customize the configuration using the IP addresses and other attributes required by your network and device.

If you need to modify a predefined FlexConfig object, copy the object, make changes to the copy, and save it with a new name. You cannot directly edit a predefined FlexConfig object.

Although you might be able to configure other ASA-based features using FlexConfig, the configuration of those features has not been tested. If an ASA feature overlaps with something that you can configure in Cloud-Delivered Firewall Management Center policies, do not attempt to configure it through FlexConfig.

For example, Snort inspection includes the HTTP protocol, so do not enable ASA-style HTTP inspection. You cannot add **http** to the enableInspectProtocolList object; this restriction prevents you from misconfiguring your device. Instead, configure the access control policy to perform application or URL filtering, as needed, to implement your HTTP inspection requirements.

Table 90: Predefined FlexConfig objects

FlexConfig Object Name	Description	Associated Text Objects
Default_Inspection_Protocol_Disable	Disables protocols in the global_policy default policy map.	disableInspectProtocolList
Default_Inspection_Protocol_Enable	Enables protocols in the global_policy default policy map.	enableInspectProtocolList
Inspect_IPv6_Configure	Configures IPv6 inspection in the global_policy policy map, logging and dropping traffic based on IPv6 header contents.	IPv6RoutingHeaderDropLogList, IPv6RoutingHeaderLogList, IPv6RoutingHeaderDropList.
Inspect_IPv6_UnConfigure	Clears and disables IPv6 inspection.	—
ISIS_Configure	Configures global parameters for IS-IS routing.	isIsNet, isIsAddressFamily, isISType

FlexConfig Object Name	Description	Associated Text Objects
ISIS_Interface_Configuration	Interface level IS-IS configuration.	isIsAddressFamily, IsIsIntfList Also uses the system variable SYS_FTD_ROUTED_INTF_MAP_LIST
ISIS_Unconfigure	Clears the IS-IS router configuration on the device.	—
ISIS_Unconfigure_All	Clears the IS-IS router configuration from the device, including the router assignment from the device interface.	—
NGFW_TCP_NORMALIZATION	Modifies the default TCP normalization configuration.	—
Policy_Based_Routing	To use this example configuration, copy it, modify the interface name, and use the r-map-object text object to identify a route map object in the object manager.	—
Policy_Based_Routing_Clear	Clears Policy Based Routing configurations from the device.	—
Sysopt_AAA_radius	Ignores the authentication key in RADIUS accounting responses.	—
Sysopt_AAA_radius_negate	Negates the Sysopt_AAA_radius configuration.	—
Sysopt_basic	Configures sysopt wait time , maximum segment size for TCP packets, and detailed traffic statistics.	tcpMssMinimum, tcpMssBytes
Sysopt_basic_negate	Clears sysopt_basic detailed traffic statistics, wait time, and TCP maximum segment size.	—
Sysopt_clear_all	Clears all sysopt configurations from the device.	—
Sysopt_noproxyarp	Configures noproxy-arp CLIs.	Uses system variable SYS_FW_NON_INLINE_INTF_NAME_LIST
Sysopt_noproxyarp_negate	Clears Sysopt_noproxyarp configurations.	Uses system variable SYS_FW_NON_INLINE_INTF_NAME_LIST
Sysopt_Preserve_Vpn_Flow	Configures syopt preserve VPN flow.	—
Sysopt_Preserve_Vpn_Flow_negate	Clears the Sysopt_Preserve_Vpn_Flow configuration.	—
Sysopt_Reclassify_Vpn	Configures sysopt reclassify vpn.	—

FlexConfig Object Name	Description	Associated Text Objects
Sysopt_Reclassify_Vpn_Negate	Negates sysopt reclassify vpn.	—
Threat_Detection_Clear	Clear the threat detection TCP Intercept configuration.	—
Threat_Detection_Configure	Configure threat detection statistics for attacks intercepted by TCP Intercept.	threat_detection_statistics
Wccp_Configure	This template provides an example for configuring WCCP.	isServiceIdentifier, serviceIdentifier, wccpPassword
Wccp_Configure_Clear	Clears WCCP configurations.	—

Deprecated FlexConfig objects

This table lists objects that configure features you can now configure natively in the GUI. Discontinue using these objects at the earliest convenience.

Table 91: Deprecated predefined FlexConfig objects

Deprecated Version	FlexConfig Object	Description	Now Configure In
7.3	DHCPv6_Prefix_Delegation_Configure	Configure one outside (Prefix Delegation client) and one inside interface (recipient of delegated prefix) for IPv6 prefix delegation. To use this template, copy it and modify the variables. Associated text objects: pdoutside, pdinside Also uses the system variable SYS_FTD_ROUTED_INTF_MAP_LIST	Interface IPv6 Settings.
7.3	DHCPv6_Prefix_Delegation_UnConfigure	Removes the DHCPv6 prefix delegation configuration.	Interface IPv6 Settings.
6.3	Default_DNS_Configure	Configure the Default DNS group, which defines the DNS servers that can be used when resolving fully-qualified domain names on the data interfaces. Associated text objects: defaultDNSNameServerList, defaultDNSParameters	Platform settings.

Deprecated Version	FlexConfig Object	Description	Now Configure In
6.3	DNS_Configure	Configure DNS servers in a non-default DNS server group. Copy the object to change the name of the group.	**DNS Server Group** in the object manager.
6.3	DNS_UnConfigure	Removes the DNS server configuration performed by Default_DNS_Configure and DNS_Configure. Copy the object to change the DNS server group names if you altered DNS_Configure.	**DNS Server Group** in the object manager.
7.2	Eigrp_Configure	Configures EIGRP routing next-hop, auto-summary, router-id, eigrp-stub. Associated text objects: eigrpAS, eigrpNetworks, eigrpDisableAutoSummary, eigrpRouterId, eigrpStubReceiveOnly, eigrpStubRedistributed, eigrpStubConnected, eigrpStubStatic, eigrpStubSummary	For all EIGRP objects, see EIGRP, on page 1479. The system does allow you to deploy post-upgrade, but also warns you to redo your EIGRP configurations. To help you with this process, we provide a command-line migration tool.
7.2	Eigrp_Interface_Configure	Configures EIGRP interface authentication mode, authentication key, hello interval, hold time, split horizon. Associated text objects: eigrpIntfList, eigrpAS, eigrpAuthKey, eigrpAuthKeyId, eigrpHelloInterval, eigrpHoldTime, eigrpDisableSplitHorizon Also uses the system variable SYS_FTD_ROUTED_INTF_MAP_LIST	
7.2	Eigrp_Unconfigure	Clears EIGRP configuration for an autonomous system from the device.	
7.2	Eigrp_Unconfigure_all	Clears all EIGRP configurations.	
7.4	Netflow_Add_Destination	Creates and configures a Netflow export destination. Associated text objects: Netflow_Destinations, netflow_Event_Types	Platform settings.
7.4	Netflow_Clear_Parameters	Restores Netflow export global default settings.	Platform settings.

Deprecated Version	FlexConfig Object	Description	Now Configure In
7.4	Netflow_Delete_Destination	Deletes a Netflow export destination. Associated text objects: Netflow_Destinations, netflow_Event_Types	Platform settings.
7.4	Netflow_Set_Parameters	Sets global parameters for Netflow export. Associated text objects: netflow_Parameters	Platform settings.
6.3	TCP_Embryonic_Conn_Limit	Configures embryonic connection limits to protect against SYN Flood Denial of Service (DoS) attacks. Associated text objects: tcp_conn_misc, tcp_conn_limit	Service policy.
6.3	TCP_Embryonic_Conn_Timeout	Configures embryonic connection timeouts to protect against SYN Flood Denial of Service (DoS) attacks. Associated text objects: tcp_conn_misc, tcp_conn_timeout	Service policy.
7.2	VxLAN_Clear_Nve	Removes the NVE 1 configured when VxLAN_Configure_Port_And_Nve is used from the device.	For all VxLAN objects, see Configure VXLAN Interfaces, on page 1062. If you configured VXLAN interfaces with FlexConfig in a previous version, they continue to work. In fact, FlexConfig takes precedence in this case—if you redo your VXLAN configurations in the web interface, remove the FlexConfig settings.
7.2	VxLAN_Clear_Nve_Only	Clears the NVE configured on the interface when deployed.	
7.2	VxLAN_Configure_Port_And_Nve	Configures VLAN port and NVE 1. Associated text objects: vxlan_Port_And_Nve	

Deprecated Version	FlexConfig Object	Description	Now Configure In
7.2	VxLAN_Make_Nve_Only	Sets an interface for NVE only. Associated text objects: vxlan_Nve_Only Also uses system variables SYS_FTD_ROUTED_MAP_LIST and SYS_FTD_SWITCHED_INTF_MAP_LIST	
7.2	VxLAN_Make_Vni	Creates a VNI interface. After deploying this you have to unregister and re-register the device to properly discover the VNI interface. Associated text objects: vxlan_Vni	

Predefined text objects

There are several predefined text objects. These objects are associated with variables used in the predefined FlexConfig objects. In most cases, you must edit these objects to add values if you use the associated FlexConfig object, or you will see errors during deployment. Although some of these objects contain default values, others are empty.

For information on editing text objects, refer to Configure FlexConfig text objects, on page 1337.

Name	Description	Associated FlexConfig Object
defaultDNSNameServerList (Deprecated.)	The DNS server IP address to configure in the Default DNS group. Starting with version 6.3, configure DNS for the data interfaces in the Threat Defense Platform Settings policy.	Default_DNS_Configure
defaultDNSParameters (Deprecated.)	The parameters to control DNS behavior for the default DNS server group. The object contains separate entries, in order, for retries, timeout, expire-entry-timer, poll-timer, domain-name. Starting with version 6.3, configure DNS for the data interfaces in the Threat Defense Platform Settings policy.	Default_DNS_Configure
disableInspectProtocolList	Disables protocols in the default policy map (global_policy).	Disable_Default_Inspection_Protocol
dnsNameServerList	The DNS server IP address to configure in a user-defined DNS group.	DNS_Configure

Name	Description	Associated FlexConfig Object
dnsParameters	The parameters to control DNS behavior for a non-default DNS server group. The object contains separate entries, in order, for retries, timeout, domain-name, name-server-interface.	DNS_Configure
enableInspectProtocolList	Enables protocols in the default policy map (global_policy). You are prevented from adding protocols whose inspection conflicts with Snort inspection.	Enable_Default_Inspection_Protocol
IPv6RoutingHeaderDropList	The list of IPv6 routing header types that you want to disallow. IPv6 inspection drops packets that contain these headers without logging the drop.	Inspect_IPv6_Configure
IPv6RoutingHeaderDropLogList	The list of IPv6 routing header types that you want to disallow and log. IPv6 inspection drops packets that contain these headers and sends a syslog message about the drop.	Inspect_IPv6_Configure
IPv6RoutingHeaderLogList	The list of IPv6 routing header types that you want to allow but log. IPv6 inspection allows packets that contain these headers, but sends a syslog message about the existence of the header.	Inspect_IPv6_Configure
isIsAddressFamily	The IPv4 or IPv6 address family.	ISIS_Configure ISIS_Interface_Configuration
IsIsIntfList	List of logical interface names.	ISIS_Interface_Configuration
isIsISType	IS Type (level-1, level-2-only or level-1-2).	ISIS_Configure
isIsNet	Network entity.	ISIS_Configure
isServiceIdentifier	When false, uses the standard **web-cache** service identifier.	Wccp_Configure
netflow_Destination	Defines a single NetFlow export destination's interface, destination, and UDP port number.	Netflow_Add_Destination
netflow_Event_Types	Defines the types of events to be exported for a destination as any subset of: **all**, **flow-create**, **flow-defined**, **flow-teardown**, **flow-update**.	Netflow_Add_Destination

Name	Description	Associated FlexConfig Object
netflow_Parameters	Provides the NetFlow export global settings: active refresh interval (number of minutes between flow update events), delay (flow create delay in seconds; default 0 = command will not appear), and template time-out rate in minutes.	Netflow_Set_Parameters
PrefixDelegationInside	Configures the inside interface for DHCPv6 prefix delegation. The object includes multiple entries, in order, interface name, IPv6 suffix with prefix length, and prefix pool name.	None, but could be used with a copy of DHCPv6_Prefix_Delegation_Configure.
PrefixDelegationOutside	Configure the outside DHCPv6 prefix delegation client. The object includes multiple entries, in order, interface name and IPv6 prefix length	None, but could be used with a copy of DHCPv6_Prefix_Delegation_Configure.
serviceIdentifier	Dynamic WCCP service identifier number.	Wccp_Configure
tcp_conn_limit (Deprecated.)	Parameters used for configuring the TCP embryonic connection limits. Starting with version 6.3, configure these features in the Threat Defense Service Policy, which you can find on the Advanced tab of the access control policy assigned to the device.	TCP_Embryonic_Conn_Limit
tcp_conn_misc (Deprecated.)	Parameters used for configuring the TCP embryonic connection settings. Starting with version 6.3, configure these features in the Threat Defense Service Policy, which you can find on the Advanced tab of the access control policy assigned to the device.	TCP_Embryonic_Conn_Limit, TCP_Embryonic_Conn_Timeout
tcp_conn_timeout (Deprecated.)	Parameters used for configuring the TCP embryonic connection timeouts. Starting with version 6.3, configure these features in the Threat Defense Service Policy, which you can find on the Advanced tab of the access control policy assigned to the device.	TCP_Embryonic_Conn_Timeout
tcpMssBytes	Maximum segment size in bytes.	Sysopt_basic
tcpMssMinimum	Checks whether to set maximum segment size (MSS), which is set only if this flag is true.	Sysopt_basic

Name	Description	Associated FlexConfig Object
threat_detection_statistics	Parameters used for threat detection statistics for TCP Intercept.	Threat_Detection_Configure
vxlan_Nve_Only	Parameters for configuring NVE-only on interface: • logical name of interface • IPv4 address (optional for routed interface) • IPv4 netmask (optional for routed interface)	VxLAN_Make_Nve_Only
vxlan_Port_And_Nve	Parameters used for configuring ports and NVE for VXLAN: • vxlan port • source interface (logical name) • type (peer or mcast) • Peer IP Address or default-mcast-group	VxLAN_Configure_Port_And_Nve
vxlan_Vni	Parameters used for creating VNI: • Interface number (1-10000) • segment-id (1-16777215) • nameif (Logical Name of the interface) • type (routed or transparent) • IP address (used in case of routed mode device) or bridge-group number (used in case of transparent mode device) • netmask (If device is in routed mode) or unused	VxLAN_Make_Vni
wccpPassword	WCCP password.	Wccp_Configure

Requirements and prerequisites for FlexConfig policies

This topic outlines the model support, domain requirements, and user role prerequisites required to configure and deploy FlexConfig policies.

Model support

Firewall Threat Defense

Supported domains

Any

User roles

Admin

Guidelines and limitations for FlexConfig

Be cautious with FlexConfig deployments, as configuration mistakes trigger rollbacks that can disrupt your network.

- If you make a mistake in the FlexConfig policy, the system will roll back all changes included in the deployment attempt that includes the failed FlexConfig. Because rollback due to a failed deployment includes clearing the configuration, this can be disruptive to your network.
- Consider timing deployments that include FlexConfig changes to non-business hours. Also, consider isolation the deployment so it includes just FlexConfig changes, and no other policy updates.

When you use the VxLAN_Make_VNI object, you must deploy the same FlexConfig to all units in a cluster or high availability pair before you form the cluster or high availability pair. The Management Center requires the VXLAN interfaces to match on all devices before forming the cluster or high availability pair.

The **webvpn keepout** command must be configured only within the global webvpn configuration mode. The webvpn group-policy attributes or username attributes configuration modes do not support the **keepout** command. Do not include the **keepout** command in a group-policy FlexConfig object.

If you configure any service that applies to connections, such as SIP inspection, go to the device CLI and enter the **clear conn** command to clear connections. When the connections are rebuilt, the new configuration is applied to the sessions.

QoS rules and service policy rules (in the access control policy) are translated to commands on the firewall. This includes creating policy maps. There can be only one policy map per interface. The system automatically combines QoS and service policy configurations in a single policy map per interface.

However, you can also create policy maps manually using FlexConfig. For lower-end firewalls, such as the Secure Firewall 3100, you can create up to 64 policy maps in the configuration. For more powerful firewalls, such as the Firepower 4100, you can create up to 128. Be careful not to exceed this limit, or deployment will fail. You can view the device configuration from the CLI or use the deployment transcript to see how rules are translated to commands.

Customize device configuration with FlexConfig policies

Use FlexConfig policies to customize the configuration of a Firewall Threat Defense device when standard configuration methods are insufficient.

Before using FlexConfig, try to configure all the policies and settings you need using the other features in Cloud-Delivered Firewall Management Center. FlexConfig is a method of last resort to configure ASA-based

features that are compatible with Firewall Threat Defense but which are not otherwise configurable in Cloud-Delivered Firewall Management Center.

Before you begin

Follow these steps to customize device configuration with FlexConfig policies:

Procedure

Step 1 Determine the CLI command sequence that you want to configure.

If you have a functioning configuration on an ASA device, use **show running-config** to get the sequence of commands that you need. Make adjustments to items such as interface names and IP addresses as needed.

If this is for a new feature, it is best to try to implement it on an ASA device in a lab setting to verify that you have the correct command sequence.

For more information, see the following topics:

- Recommended usage for FlexConfig policies, on page 1310
- CLI commands in FlexConfig objects, on page 1310

Step 2 Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

Examine the predefined FlexConfig objects to determine if any will be able to generate the commands you need. Click **View** () to see the object contents. If an existing object is close to what you want, start by making a copy of the object, and then edit the copy. See Predefined FlexConfig objects, on page 1322.

Examining the objects will also give you an idea of the structure, command syntax, and expected sequencing for a FlexConfig object.

Note
If you find any objects that you will use, either directly or as copies, examine the Variables list at the bottom of the object. Make note of the variable names, except those in all capitals that start with SYS, which are system variables. These variables are text objects that you will probably need to edit and define overrides for, especially if the default value column shows the object has no value.

Step 3 If you need to create your own FlexConfig objects, determine what variables you will need and create the associated objects.

The CLI you need to deploy might contain IP addresses, interface names, port numbers, and other parameters that you might want to adjust over time. These are best isolated into variables, which point to objects that contain the necessary values. You might also need variables for strings that are part of the configuration and could change over time.

Also, determine if you need different values for each device to which you will assign the policy. For example, you might want to configure the feature on three devices, but you might need to specify a different interface name or IP address on a given command for each of these devices. If you need to customize the object for each device, ensure that you enable overrides when creating the object and then define the override values per device.

See the list of topics for an explanation of the various types of variables and how to configure the related objects when necessary.

- FlexConfig variables, on page 1314

- FlexConfig policy object variables, on page 1320
- FlexConfig system variables, on page 1321
- Configure FlexConfig text objects, on page 1337

Step 4 If you are using the predefined FlexConfig objects, edit the text objects used as variables.

See Configure FlexConfig text objects, on page 1337.

Step 5 (If necessary) Configure FlexConfig objects, on page 1333.

You need to create objects only if the predefined objects cannot do the job.

Step 6 Configure the FlexConfig policy, on page 1339.

a) Set target devices for a FlexConfig policy, on page 1340.

You can also assign the policy to devices when you create the policy. The policy must have at least one assigned device before you can preview it.

b) Preview the FlexConfig policy, on page 1341.

You must save changes before you can preview the policy.

Verify that the generated commands are the ones intended, and that all variables are resolving correctly.

Step 7 Choose **Deploy** > **Deployment** in the menu bar.

a) Select the devices assigned to the policy, and click **Deploy**.

Wait for deployment to complete.

Step 8 Verify the deployed configuration, on page 1342.

Step 9 (If necessary) Remove features configured using FlexConfig, on page 1344.

Unlike other types of policy, simply unassigning a FlexConfig from a device might not remove the related configuration. If you want to remove a FlexConfig-generated configuration, follow the cited procedure.

If you are removing a feature because it is now directly supported by the product, check Convert from FlexConfig to managed feature, on page 1345.

Configure FlexConfig objects

Before you begin

Keep the following in mind:

- FlexConfig objects translate into commands that are then deployed to the device. These commands are already issued in global configuration mode. Therefore, do not include the **enable** and **configure terminal** commands as part of the FlexConfig object.
- Determine what types of variables you will need, and create any policy objects that you will require. You cannot create objects for variables while editing a FlexConfig object.
- Ensure that your commands do not conflict in any way with the VPN or access control configuration on the devices.

- If there is more than one set of commands for an interface, only the last set of commands is deployed. Therefore, we recommend you not use beginning and ending commands to configure interfaces. For an example of configuring interfaces, refer to the ISIS_Interface_Configuration predefined FlexConfig object.

There are several predefined FlexConfig objects that you can use directly, or you can make copies if you need to edit them. You can also create your own objects from scratch. A FlexConfig object's content can range from a single simple command string to elaborate CLI command structures that use variables and scripting commands to deploy commands whose content can differ from device to device or deployment to deployment.

You can also create FlexConfig policy objects when defining FlexConfig policies.

Procedure

Step 1 Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object** .

Step 2 Do one of the following:

- Click **Add FlexConfig Object** to create a new object.
- Click **Edit** ( ) to edit an existing object.
- Click **View** ( ) to see the contents of a predefined object.
- If you want to edit a predefined object, click **Clone** ( ) to create a new object with the same contents.

Step 3 Enter a **Name** and optionally, a description for the object.

Step 4 In the object body area, enter the commands and instructions to produce the required configuration.

The object content is a sequence of scripting commands and configuration commands that generate a valid ASA software command sequence. The Firewall Threat Defense device uses ASA software commands to configure some features. For more information on scripting and configuration commands, refer to:

- Template scripts, on page 1314
- CLI commands in FlexConfig objects, on page 1310

You can use variables to supply information that can be known only at runtime, or which can differ from device to device. You simply type in processing variables, but you must use the **Insert** menu to add variables that are associated with policy objects or system variables, or which are secret keys. For a complete discussion of variables, refer to FlexConfig variables, on page 1314 .

- To insert system variables, choose **Insert** > **Insert System Variable** > ***Variable Name*** . For a detailed explanation of these variables, see FlexConfig system variables, on page 1321 .
- To insert policy object variables, choose **Insert** > **Insert Policy Object** > ***Object Type*** , selecting the appropriate type of object. Then, give the variable a name (which can be the same name as the associated policy object), select the object to associate with the variable, and click **Save** . For a detailed explanation of these types, see FlexConfig policy object variables, on page 1320 . For more detail on the procedure, see Add a policy object variable to a FlexConfig object, on page 1336 .
- To insert secret key variables, choose **Insert** > **Secret Key** and define the variable name and value. For more detail on the procedure, see Configure secret keys, on page 1337 .

Note
You must use the **Insert** menu to create a new policy object or system variable. However, for subsequent uses of that variable, you must type it in, $ included. This is also true for system variables: the first time you use it, add it from the **Insert** menu. Then, type it out for subsequent uses. If you use the **Insert** menu more than once for a system variable, the system variable is added to the Variables list multiple times, and the FlexConfig will not validate, meaning you cannot save your changes. For processing variables (those not associated with a policy object or system variable), simply type in the variable. If you are adding a secret key, always use the **Insert** menu. Secret key variables do not show up in the Variables list.

Step 5 Choose the deployment frequency and type.

- **Deployment** —Whether to deploy the commands in the object **Once** or **Everytime** . The only way to choose the right option is to test the results of deployment.

  Start by selecting **Everytime** . Then, after you attach the object to a FlexConfig policy, deploy the configuration. After a successful deployment, come back to the FlexConfig policy and preview the configuration for one of the assigned devices as described in Preview the FlexConfig policy, on page 1341 . If the section labeled `###CLI generated from managed features ###` contains commands that clear or negate the commands in the object, and the `###Flex-config Appended CLI ###` section contains the commands to reconfigure the feature, you know that **Everytime** is the right option.

  Even if you do not see negate commands, make some minor change to the device configuration, then run another deployment. If the deployment completes successfully, you can check the deployment transcript (see Verify the deployed configuration, on page 1342). If you see that the commands were issued again (even when they were already configured) without error, then you can keep **Everytime** .

  Change to **Once** only if the system does not first negate the commands in the object before issuing them again, or if the deployment results in errors that are specific to the command. In some cases, the system does not allow you to issue a command that is already configured, but this is the exception.

  Some additional tips:

  - If the FlexConfig object points to system-managed objects such as network or ACL objects, choose **Everytime** . Otherwise, updates to the objects might not get deployed.
  - Use **Once** if the only thing you do in the object is to clear a configuration. Then, remove the object from the FlexConfig policy after the next deployment.

- **Type** —Select one of the following:
  - **Append** —(The default.) Commands in the object are put at the end of the configurations generated from the Cloud-Delivered Firewall Management Center policies. You must use Append if you use policy object variables, which point to objects generated from managed objects. If commands generated for other policies overlap with those specified in the object, you should select this option so your commands are not overwritten. This is the safest option.
  - **Prepend** —Commands in the object are put at the beginning of the configurations generated from the Cloud-Delivered Firewall Management Center policies. You would typically use prepend for commands that clear or negate a configuration.

Step 6 (Optional.) Click **Validate** () above the object body to check the integrity of the script.

The object is always validated when you click **Save**. You cannot save an invalid object.

Step 7 Click **Save** .

What to do next

- If an active policy references your object, deploy configuration changes.

Add a policy object variable to a FlexConfig object

Insert variables into a FlexConfig policy object that are associated with other types of policy object, allowing the variables to resolve to the names or content of the associated object when the FlexConfig is deployed to a device.

You can insert variables into a FlexConfig policy object that are associated with other types of policy object. When the FlexConfig is deployed to a device, these variables resolve to the names or content of the associated object.

Use the following procedure for the first use of a policy object variable in a FlexConfig object. If you need to refer to the object again, type in the variable (including the $ sign). To understand how to use these variables, see Variable processing, on page 1315.

Before you begin

For information on editing a FlexConfig Object, see Configure FlexConfig objects, on page 1333.

Procedure

Step 1 While editing a FlexConfig Policy Object, choose **Insert** > **Insert Policy Object** > ***Object Type***, selecting the appropriate type of object.

Step 2 Enter a name for the variable, and optionally, a description.

The name must be unique within the context of the FlexConfig object. It cannot include spaces. You are allowed to use the exact same name as the object associated with the variable.

Step 3 Select the object to associate with the variable and click **Add** to move it to the **Selected Object** list.

You can associate a variable with a single object only.

Note
For text objects, you can select any of the predefined objects as needed. However, many of these objects have no default values. You must update the objects to add the required values either directly or as overrides for the device to which you will deploy the FlexConfig object. Trying to deploy a FlexConfig without updating these objects typically results in deployment errors.

Step 4 Click **Save**.

The variable appears in the Variables list at the bottom of the FlexConfig object editor.

Configure secret keys

Secret keys allow you to mask sensitive information such as passwords in FlexConfig objects, preventing the dissemination of sensitive information while providing special treatment for these variables.

A secret key is any single-string variable whose content you want to mask, such as passwords. The system provides special treatment for these variables to help you prevent the dissemination of sensitive information.

Secret key variables do not show up in the Variables list in the FlexConfig object.

Unlike other types of variables, you can use the **Insert** command every time you need to insert a given secret key variable. With respect to processing, these variables behave like single-value text object variables; see Single value variables, on page 1315.

Note Any data defined in a secret key variable is masked from users except when previewing a FlexConfig policy. In addition, if you export a FlexConfig policy, the content of any secret key variable is erased. When you import the policy, you will need to manually edit each secret key variable to enter the data.

Before you begin

For information on editing a FlexConfig Object, see Configure FlexConfig objects, on page 1333.

Procedure

Step 1 While editing a FlexConfig Policy Object, choose **Insert** > **Secret Key**.

Step 2 In the **Insert Secret Key** dialog box, do any of the following:

- To create a new key, click **Add Secret Key**, then fill in the following information and click **Add**.
  - **Secret Key Name**—The name of the variable. This name appears in the FlexConfig object prefixed with @.
  - **Password**, **Confirm Password**—The secret string, which is masked with asterisks as you type.
- To insert a secret key variable in the FlexConfig object, select the check box for the variable.
- To edit the value of a secret key variable, click **Edit** ( ) for the variable. Make your changes and click **Add**.
- To delete a secret key variable, click **Delete** ( ) for the variable.

Step 3 Click **Save**.

Configure FlexConfig text objects

Use text objects in FlexConfig objects as the target of policy object variables. During deployment, variables that point to text objects are replaced by the content of the text object.

Use text objects in FlexConfig objects as the target of policy object variables. You can use variables to supply information that can be known only at runtime, or which can differ from device to device. During deployment, variables that point to text objects are replaced by the content of the text object.

Text objects contain free-form strings, which can be keywords, interface names, numbers, IP addresses, and so forth. The content depends on how you will use the information within a FlexConfig script.

Before creating or editing a text object, determine exactly what content you will need. This includes how you intend to process the object, which will help you decide between creating a single string or multiple string object. Read the following topics:

- FlexConfig variables, on page 1314
- Variable processing, on page 1315

Procedure

Step 1 Choose **Objects** > **Object Management** > **FlexConfig** > **Text Object**.

Step 2 Do one of the following:

- Click **Add Text Object** to create a new object.
- Click **Edit** ( ) to edit an existing object. You are allowed to edit the predefined text objects, which is required if you intended to use the predefined FlexConfig objects.

Step 3 Enter a **Name** and optionally, a description for the object.

Step 4 (New objects only.) Choose a **Variable Type** from the drop-down list:

- **Single**—If the object should contain a single text string.
- **Multiple**—If the object should contain a list of text strings.

You cannot change the variable type after you save the object.

Step 5 If the variable type is **Multiple**, use the up and down arrows to specify a **Count**.

Rows are added or removed from the object as you change the number.

Step 6 Add content to the object.

You can either click in the text box next to a variable number and type in a value, or you can set up device overrides for each device that will be assigned a FlexConfig object that uses the text object. You can also do both, in which case the values configured in the base object act as default values in cases where an override does not exist for a given device.

When editing predefined objects, it is a good practice to use device overrides, so that the system defaults remain in place for other users who might need to use the object in different FlexConfig policies. The approach you take depends on the requirements of your organization.

Tip
Some predefined objects require multiple values where each value serves a specific purpose. Read the description text carefully to determine the expected values in the object. In some cases, the instructions specify that you must use overrides instead of changing the base values. In the case of enableInspectProtocolList, you are prevented from entering protocols whose inspection is incompatible with Snort inspection.

If you decide to use device overrides, do the following.

a) Check the check box of **Allow Overrides**.
b) Expand the Overrides area (if necessary) and click **Add**.

   If an override already exists for the device, click edit for the override to change it.

c) On **Targets** in the Add Object Override dialog box, select the device for which you are defining values and click **Add** to move it to the Selected Devices list.
d) Click **Override**, adjust the **Count** as needed, then click in the variable fields and type in the values for the device.
e) Click **Add**.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Configure the FlexConfig policy

Configure a FlexConfig policy to manage device configurations using shared policies that can be assigned to multiple devices.

A FlexConfig policy contains two ordered lists of FlexConfig objects, one prepended list and one appended list. For an explanation of prepend/append, see Configure FlexConfig objects, on page 1333.

FlexConfig policies are shared policies that you can assign to multiple devices.

Procedure

Step 1 Choose **Devices** > **FlexConfig**.

Step 2 Do one of the following:

- Click **New Policy** to create a new FlexConfig Policy. You are prompted to enter a name. Optionally, select devices in the Available Devices list and click **Add to Policy** to assign devices. Click **Save**.
- Click **Edit** () to edit an existing Policy. You can change the name or description by clicking them in edit mode.
- Click **Copy** () to create a new policy with the same contents. You are prompted for a name. Device assignments are not retained for the copy.
- Click delete to remove a policy you no longer need.

Step 3 Select the FlexConfig objects required for the policy from the **Available FlexConfig** list and click > to add them to the policy.

Objects are automatically added to the prepended or appended list based on the deployment type specified in the FlexConfig object.

To remove a selected object, click **Delete** ( ) next to an object.

Step 4 For each selected object, click **View** ( ) next to the object to identify the variables used in the object.

Except for system variables, which start with SYS, you need to ensure that the objects associated with the variables are not empty. A blank or brackets with nothing between them, [ ], indicate an empty object. You will need to edit these objects before deploying the policy.

Note
If you use object overrides, those values will not show up in this view. Thus, an empty default value does not necessarily mean that you have not updated the object with the required values. Previewing the configuration will show whether the variables resolve correctly for a given device. See Preview the FlexConfig policy, on page 1341.

Step 5 Click **Save**.

What to do next

- Set target devices for the policy; see Set target devices for a FlexConfig policy, on page 1340.
- Deploy configuration changes.

Set target devices for a FlexConfig policy

This task allows you to assign or change device assignments for a FlexConfig policy after creation.

When you create a FlexConfig policy, you can select the devices that use the policy. You can subsequently change device assignments for the policy as described below.

Note Normally, when you unassign a policy from a device, the system automatically removes the associated configuration upon the next deployment. However, because FlexConfig objects are scripts for deploying customized commands, simply unassigning a FlexConfig policy from a device does not remove the commands that were configuring by the FlexConfig objects. If your intention is to remove FlexConfig-generated commands from a device's configuration, refer to Remove features configured using FlexConfig, on page 1344.

Procedure

Step 1 Choose **Devices** > **FlexConfig** and edit a FlexConfig policy.

Step 2 Click **Policy Assignments**.

Step 3 On **Targeted Devices**, build your target list:

- Add—Choose one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**. You can assign the policy to devices, high availability pairs, and clustered devices.
- Delete—Click **Delete** ( ) next to a single device, or select multiple devices, right-click, then choose **Delete Selection**.

Step 4 Click **OK** to save your selection.

Step 5 Click **Save** to save the FlexConfig policy.

What to do next

- Deploy configuration changes.

Preview the FlexConfig policy

The FlexConfig policy preview allows you to validate and troubleshoot your configuration before deployment. It translates FlexConfig objects into CLI commands and resolves variables based on device-specific configurations.

Preview a FlexConfig policy to see how the FlexConfig objects get translated into CLI commands. The preview shows the commands that will be generated for a selected device from the scripts and variables used in the FlexConfig objects. The variables are resolved based on the configuration for the device, so you get a clear idea of what will be deployed.

Use the preview to look for potential problems in the FlexConfig objects. Correct the objects until the preview shows the expected results.

You must preview the configuration separately for each device, because the variables can resolve differently based on the device configuration.

Procedure

Step 1 Choose **Devices** > **FlexConfig** and edit a FlexConfig policy.

Step 2 If there are any pending changes, click **Save**.

The preview shows results only for those FlexConfig objects that were in the most recently saved version of the policy. You must save the policy to see a preview of newly-added objects.

Step 3 Click **Preview Config**.

Step 4 Choose a device from the **Select Device** drop-down list.

The system retrieves information from the device and configured policies, and determines what CLI commands will be generated on the next deployment to the device. You can select the output and use Ctrl+C to copy it to the clipboard, where you can paste it into a text file for further analysis.

The preview includes the following sections:

- Flex-config Prepended CLI—These are the commands generated by FlexConfigs that are prepended to the configuration.
- CLI generated from managed features—These are commands generated for policies configured in the Cloud-Delivered Firewall Management Center. Commands are generated for new or changed policies since the last successful deployment to the device. These commands do not represent all commands needed to implement the assigned policies. No commands in this section are generated from FlexConfig objects.

- Flex-config Appended CLI—These are the commands generated by FlexConfigs that are appended to the configuration.

Step 5 Click **Close** to close the preview dialog.

Verify the deployed configuration

This task helps you confirm that your FlexConfig policy deployment completed successfully and that the device configuration matches your expectations.

After you deploy a FlexConfig policy to a device, verify that the deployment was successful and that the resulting configuration is what you expected. Also, verify that the device is performing as expected.

Procedure

Step 1 To verify that deployment was successful:

Check the deployment status and review any error messages or warnings.

- Click **Notifications** in the menu bar, which is unnamed between **Deploy** and **System**.

  The icon looks like one of the following, and it might include a number if there are errors:

  - **Indicates No Warnings** — Indicates no warnings or errors are present on the system.
  - Indicates **One or More Warnings** — Indicates one or more warnings and no errors are present on the system.
  - **Indicates One or More Errors** — Indicates one or more errors and any number of warnings are present on the system.

- On **Deployments**, verify that the deployment was successful.
- To see more detailed information, especially for failed deployments, click **Show History**.
- Select the deployment job in the list of jobs in the left column.

  Jobs are listed in reverse chronological order, with the most recent job at the top of the list.

- Click download in the **Transcript** column for the device in the right column.

  The deployment transcript includes commands sent to the device, and any responses returned from the device. These response can be informative messages or error messages. For failed deployments, look for messages that indicate errors with the commands that you sent through FlexConfig. These errors can help you correct the script in the FlexConfig object that is trying to configure the commands.

  Note
  There is no distinction made in the transcript between commands sent for managed features and those generated from FlexConfig policies.

  For example, the following sequence shows that Cloud-Delivered Firewall Management Center sent commands to configure GigabitEthernet0/0 with the logical name outside. The device responded that it automatically set the security level to 0. Firewall Threat Defense does not use the security level for anything. Messages relevant to FlexConfig are in the CLI Apply section of the transcript.

```
========= CLI APPLY =========

FMC >> interface GigabitEthernet0/0
FMC >>  nameif outside
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

Step 2 Verify that the deployed configuration includes the expected commands.

You can do this by making an SSH connection to the device's management IP address. Use the **show running-config** command to view the configuration.

Alternatively, use the CLI tool within Cloud-Delivered Firewall Management Center.

- Choose **System** (⚙) > **Health** > **Monitor** and click the name of the device.

  You might need to click the open/close arrow in the **Count** column in the Status table to see any devices.

- Click **Advanced Troubleshooting**.
- Click **Threat Defense CLI**.
- Select **show** as the command, and type **running-config** as the parameter.
- Click **Execute**.

  The running configuration appears in the text box. You can select the configuration and press Ctrl+C, then paste it into a text file for later analysis.

Step 3 Verify that the device is performing as expected.

Use the **show** commands related to the feature to see detailed information and statistics. For example, if you enabled additional protocol inspections, the **show service-policy** command provides this information. The exact commands to use are feature-dependent and should be mentioned in the ASA configuration guide and command reference you used to learn how to configure the feature.

If commands that show statistics indicate that numbers are not changing (for example, hit counts, connection counts, and so forth), the configuration might be valid but not meaningful. If you know that traffic is going through the device that should show up in statistics, look for what is missing in your configuration. For example, NAT or access rules might be dropping or changing traffic before a feature can act on it.

You can use the **show** commands from an SSH session or through the Cloud-Delivered Firewall Management Center CLI tool.

However, if the **show** command that you need to use is not available directly within the Firewall Threat Defense CLI, you will need make an SSH connection to the device to use the commands. From the CLI, enter the following command sequence to enter Privileged EXEC mode within the diagnostic CLI. From there, you should be able to enter these otherwise unsupported **show** commands.

```
> system support diagnostic-cli
Attaching to Diagnostic CLI ... Press 'Ctrl+a then d' to detach.
Type help or '?' for a list of available commands.
firepower> enable
Password: <press enter, do not enter a password>
firepower#
```

Remove features configured using FlexConfig

This task allows you to manually remove configuration commands that were deployed using FlexConfig when unassigning the FlexConfig policy from a device might not remove all of the configuration.

If you decide you need to remove a set of configuration commands you configured using FlexConfig, you might need to manually remove that configuration. Unassigning the FlexConfig policy from a device might not remove all of the configuration.

To manually remove the configuration, you create new FlexConfig objects to clear or negate the configuration commands.

Before you begin

To determine if you need to manually remove some or all of the configuration generated by an object:

1. Examine the configuration preview, as described in Preview the FlexConfig policy, on page 1341. If the **###CLI generated from managed features ###** section contains the clear or negate commands to remove all of the commands in the FlexConfig object, then you can simply remove the object from the FlexConfig policy, save, and redeploy.
2. Remove the object from the FlexConfig policy, save the change, then preview the configuration again. If the **###CLI generated from managed features ###** section still does not include the required clear or negate commands, you must follow this procedure to manually remove the configuration.

Follow these steps to remove features configured using FlexConfig:

Procedure

Step 1 Choose **Objects** > **Object Management** and create the FlexConfig Objects to clear or negate the configuration commands.

If a feature has a **clear** command that can remove all configuration settings, then use that command. For example, the predefined ISIS_Unconfigure_All object contains a single command that removes all ISIS-related configuration commands:

```
clear configure router isis
```

If there is not a **clear** command for the feature, you need to use the **no** form of each command you want to remove. For example, the predefined Sysopt_basic_negate object removes the commands configured through the predefined Sysopt_basic object.

```
no sysopt traffic detailed-statistics
```

```
no sysopt connection timewait
```

You would typically configure a FlexConfig object that removes configurations as a prepended, deploy once object.

Step 2 Choose **Devices** > **FlexConfig** and create a new FlexConfig policy or edit the existing policy.

If you want to preserve the FlexConfig policy that deploys the configuration commands, create a new policy specifically for negating the commands, and assign the devices to the policy. Then, add the new FlexConfig objects to the policy.

If you want to completely remove the FlexConfig configuration objects from all devices, you can simply delete those commands from the existing FlexConfig policy and replace them with the objects that negate the configuration.

Step 3 Click **Save** to save the FlexConfig policy.

Step 4 Click **Preview Config** and verify that the clear and negation commands are generating correctly.

Step 5 Choose **Deploy** > **Deployment** in the menu bar, select the device, and click **Deploy**.

Wait for deployment to complete.

Step 6 Verify that the commands were removed.

View the running configuration on the device to confirm that the commands are removed. For more detailed information, see Verify the deployed configuration, on page 1342.

Step 7 While editing the FlexConfig policy, click **Policy Assignments** and remove the device. Optionally, remove the FlexConfig Objects from the policy.

Assuming that the FlexConfig policy simply removes the unwanted configuration commands, there is no need to keep the policy assigned to the device after the removal is complete.

However, if the FlexConfig policy retains options that you still want configured on the device, remove the negation objects from the policy. They are no longer needed.

Convert from FlexConfig to managed feature

This procedure converts deprecated FlexConfig commands to newly supported managed features after a software upgrade.

Each software release adds managed features to the product. These are features you configure directly through policies controlled outside of FlexConfig. This can deprecate FlexConfig commands that you are currently using. Your configurations are not automatically converted. After the upgrade, you cannot assign or create FlexConfig objects using the newly deprecated commands. After upgrading software, examine your FlexConfig policies and objects.

If a feature configured through FlexConfig becomes available as a managed feature, convert the configuration to the managed feature instead of using FlexConfig. In most cases, your existing FlexConfig configurations continue to work after the upgrade, and you can still deploy. However, using deprecated commands may cause deployment issues in certain situations. Configuring a feature in both the GUI and FlexConfig is not supported.

Note Use the migration tool instead of this procedure if the tool supports the feature configuration you are migrating.

Procedure

Step 1 Remove the FlexConfig, as explained in Remove features configured using FlexConfig, on page 1344.

Step 2 Configure the settings in the newly supported managed feature.

The release notes have a list of new features for the release.

Examples for FlexConfig

This topic provides a collection of practical examples for implementing FlexConfig policies.

Configure precision time protocol (ISA 3000)

The Precision Time Protocol (PTP) is a time-synchronization protocol developed to synchronize the clocks of various devices in a packet-based network. These device clocks are generally of varying precision and stability. The protocol is designed specifically for industrial, networked measurement and control systems, and is optimal for use in distributed systems because it requires minimal bandwidth and little processing overhead.

A PTP system is a distributed, networked system consisting of a combination of PTP and non-PTP devices. PTP devices include ordinary clocks, boundary clocks and transparent clocks. Non-PTP devices include network switches, routers and other infrastructure devices.

You can configure the Firewall Threat Defense device to be a transparent clock. The Firewall Threat Defense device does not synchronize its clock with the PTP clocks. The Firewall Threat Defense device will use the PTP default profile, as defined on the PTP clocks.

When you configure the PTP devices, you define a domain number for the devices that are meant to function together. Thus, you can configure multiple PTP domains, and then configure each non-PTP device to use the PTP clocks for one specific domain.

Before you begin

Determine the domain number configured on the PTP clocks that the device should use. This example assumes the PTP domain number is 10. Also, determine the interfaces through which the system can reach the PTP clocks in the domain.

These are the guidelines for configuring PTP.

- This feature is only available on the Cisco ISA 3000 appliance.
- Cisco PTP supports multicast PTP messages only.
- PTP is available only for IPv4 networks, not for IPv6 networks.
- PTP configuration is supported on physical Ethernet data interfaces, whether stand-alone or bridge group members. It is not supported on the management interface, subinterfaces, EtherChannels, Bridge Virtual Interfaces (BVI), or any other virtual interfaces.

- PTP flows on VLAN subinterfaces are supported, assuming the appropriate PTP configuration is present on the parent interface.
- You must ensure that PTP packets are allowed to flow through the device. PTP traffic is identified by UDP destination ports 319 and 320, and destination IP address 224.0.1.129, so any access control rule that allows this traffic should work.
- In Routed firewall mode, you must enable multicast routing for PTP multicast groups. In addition, if an interface on which you enable PTP is **not** in a bridge group, you must configure the interface to join the IGMP multicast group 224.0.1.129. If the physical interface is a bridge group member, you do not configure it to join the IGMP multicast group.

Follow these steps to configure Precision Time Protocol on an ISA 3000 device:

Procedure

Step 1 (Routed mode only) Enable Multicast routing, and configure the IGMP group for the interfaces.

In Routed mode, you must enable Multicast routing. In addition, for stand-alone physical interfaces, that is, those that are not bridge group members, you must also configure the interface to join the 224.0.1.129 IGMP group. You cannot configure bridge group members to join an IGMP group, but PTP configuration on bridge group members will work without the IGMP join.

Perform this procedure for each device on which you will configure PTP.

Note
Write down the hardware names of each PTP-clock-facing interface on each device, for example, GigabitEthernet1/1.

a) Choose **Devices** > **Device Management**, and edit the device.
b) Click **Routing**.
c) Choose **Multicast Routing** > **IGMP**.
d) Check the **Enable Multicast Routing** check box.
e) Click **Join Group**.
f) Click **Add**, and in the **Add IGMP Join Group Parameters** dialog box, configure the following options and click **OK**.
   - **Interface**—Select the PTP-clock-facing stand-alone interface.
   - **Join Group**—Click **+** to add a new network object. Create a Host object with the address 224.0.1.129. When configuring additional interfaces, simply select this object. (See Create a network object, on page 2985.)

   Repeat this step for each PTP-clock-facing stand-alone interface on the device.
g) Click **Save** on the Routing page.

Step 2 Create the FlexConfig object to enable PTP globally and on the interface.

This procedure assumes that the PTP-clock-facing interface is the same on every device you are configuring. If you have used different interfaces on different devices, you need to create separate objects for each distinct combination. For example, if you use GigabitEthernet1/1 on devices A and B, GigabitEthernet1/2 on devices

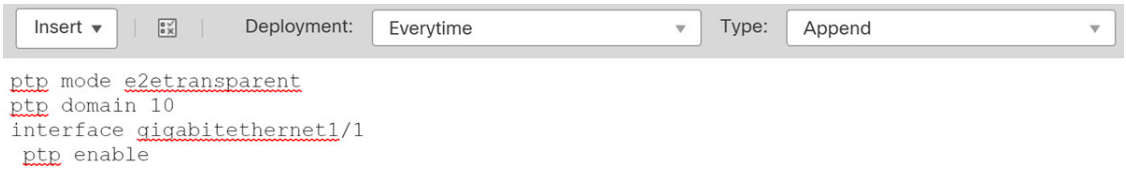
C and D, and both GigabitEthernet1/1 and 1/2 on devices E and F, you need 3 separate FlexConfig objects, and subsequently, 3 separate FlexConfig policies (explained in the next step).

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.
b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Enable_PTP.
- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features. This ensures that any other changes you make to interface configuration are configured before these commands.
- **Object body**—In the object body, type the commands needed to configure PTP globally and on each PTP-clock-facing interface. For example, the commands needed for the global configuration for PTP domain 10 and the interface configuration on GigabitEthernet1/1 are:

```
ptp mode e2etransparent
ptp domain 10
interface gigabitethernet1/1
 ptp enable
```

The object body should look similar to the following:



Step 3 Create the FlexConfig policy and assign it to the devices.

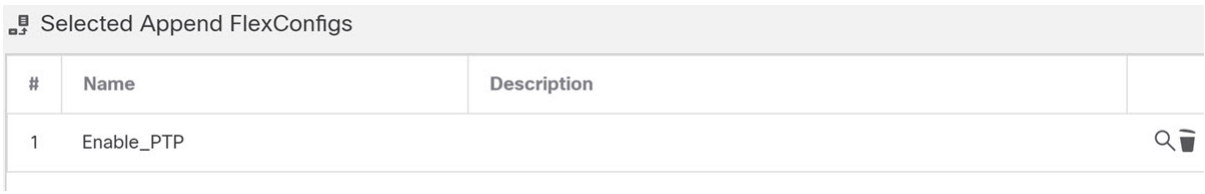
If you created multiple FlexConfig objects for different combinations of PTP-clock-facing interfaces, you need to create separate FlexConfig policies for each object, and assign those policies to the correct devices based on the interfaces you need to configure. Repeat this procedure for each group of devices.

a) Choose **Devices** > **FlexConfig**.
b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the PTP FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.



d) Click **Save**.
e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.
f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the PTP FlexConfig object look correct. These will be shown at the end of the preview. Note that you will also see commands generated from other changes you have made to managed features. For the PTP commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
ptp mode e2etransparent
ptp domain 10
interface gigabitethernet1/1
 ptp enable
```

Step 4 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the deployed configuration, on page 1342.

Step 5 Verify the PTP configuration on each device.

From an SSH or Console session into each device, verify the PTP settings:

```
> show ptp clock
 PTP CLOCK INFO
  PTP Device Type: End to End Transparent Clock
  Operation mode: One Step
  Clock Identity: 34:62:88:FF:FE:1:73:81
  Clock Domain: 10
  Number of PTP ports: 4
> show ptp port
 PTP PORT DATASET: GigabitEthernet1/1
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 1
  PTP version: 2
  Port state: Enabled

 PTP PORT DATASET: GigabitEthernet1/2
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 2
  PTP version: 2
  Port state: Disabled

 PTP PORT DATASET: GigabitEthernet1/3
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 3
  PTP version: 2
  Port state: Disabled

 PTP PORT DATASET: GigabitEthernet1/4
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 4
```

```
PTP version: 2
Port state: Disabled
```

Configure automatic hardware bypass for power failure (ISA 3000)

Configure automatic hardware bypass to ensure network traffic continues to flow between interface pairs during power outages on ISA 3000 devices.

You can enable hardware bypass so that traffic continues to flow between an interface pair during a power outage. Supported interface pairs are copper interfaces GigabitEthernet 1/1 and 1/2; and GigabitEthernet 1/3 and 1/4. If you have a fiber Ethernet model, only the copper Ethernet pair (GigabitEthernet 1/1 and 1/2) supports hardware bypass.

When hardware bypass is active, traffic passes between these interface pairs at layer 1. The Firewall Threat Defense CLI will see the interfaces as being down. No firewall functions are in place, so make sure you understand the risks of allowing traffic to pass through the device.

In CLI Console or an SSH session, use the **show hardware-bypass** command to monitor the operational status.

Before you begin

For hardware bypass to work:

- You must place the interface pairs in the same bridge group.
- You must attach the interfaces to access ports on the switch. Do not attach them to trunk ports.

We recommend that you disable TCP sequence number randomization globally using the Threat Defense Service Policy attached to the access control policy assigned to the device. By default, the ISA 3000 rewrites the initial sequence number (ISN) of TCP connections passing through it to a random number. When hardware bypass is activated, the ISA 3000 is no longer in the data path and does not translate the sequence numbers. The receiving client receives an unexpected sequence number and drops the connection, so the TCP session needs to be re-established. Even with TCP sequence number randomization disabled, some TCP connections will have to be re-established because of the link that is temporarily down during the switchover.

Follow these steps to configure automatic hardware bypass for power failure:

Procedure

Step 1 Create the FlexConfig object to enable automatic bypass.

a. Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

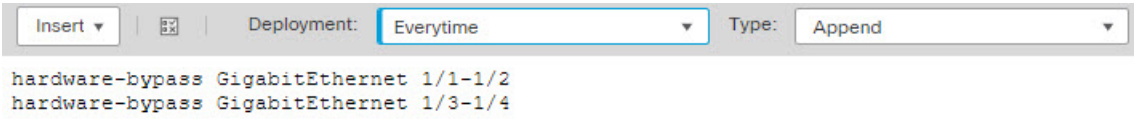
b. Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Enable_HW-Bypass.
- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features.

- **Object body**—In the object body, type the commands needed to enable automatic hardware bypass. For example, the commands needed for both possible interface pairs:

```
hardware-bypass GigabitEthernet 1/1-1/2
hardware-bypass GigabitEthernet 1/3-1/4
```

The object body should look similar to the following:



Step 2 Create the FlexConfig policy and assign it to the devices.

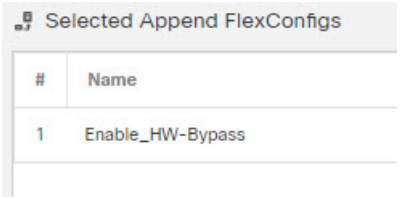
a. Choose **Devices** > **FlexConfig**.

b. Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c. Select the hardware bypass FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.



d. Click **Save**.

e. If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f. Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the hardware bypass FlexConfig object look correct. These will be shown at the end of the preview. Note that you will also see commands generated from other changes you have made to managed features. For the hardware bypass commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
hardware-bypass GigabitEthernet 1/1-1/2
hardware-bypass GigabitEthernet 1/3-1/4
```

Step 3 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the deployed configuration, on page 1342.

What to do next

If you want to manually invoke hardware bypass or manually turn it off, you need to create two FlexConfig objects:

- One that manually starts bypass, which would contain one or both of the following commands, depending on whether you want to invoke bypass for both pairs:

```
hardware-bypass manual GigabitEthernet 1/1-1/2
hardware-bypass manual GigabitEthernet 1/3-1/4
```

- One that manually turns off bypass, which would contain one or both of the following commands:

```
no hardware-bypass manual GigabitEthernet 1/1-1/2
no hardware-bypass manual GigabitEthernet 1/3-1/4
```

You would then need to add one or the other object to the FlexConfig policy, and deploy changes, to either turn bypass on or off. You would also need to immediately remove the object from the FlexConfig policy after deployment. If you manually invoke bypass, you would then need to repeat the process to turn it off again. Thus, using this manual method requires frequent and careful editing of the FlexConfig policy and additional deployments.

Migrate FlexConfig policies

This task migrates ECMP, VXLAN, and EIGRP policies from FlexConfig objects to the management center UI to enable direct policy management.

Attention This section on migrating FlexConfig policies is applicable only for migrating ECMP, VXLAN, and EIGRP policies.

The ECMP, VXLAN, and EIGRP policies were configured using the FlexConfig objects and policies in earlier versions of the Cloud-Delivered Firewall Management Center. You can now directly configure these policies in the management center UI. When you are upgrading the management center from earlier versions, the FlexConfig configuration is retained. However, to manage the policies from the UI, you must redo the configuration in the corresponding **Device (Edit)** > **Routing** page and remove the configuration from FlexConfig. To automate creation of the policies in the UI, Cloud-Delivered Firewall Management Center provides an option to migrate the policies from FlexConfig to the UI. However, it does not remove the migrated policies from the FlexConfig. For post-migration procedure, refer to Step 7, on page 1353.

Before you begin

- Ensure that the deployed FlexConfig policy is up-to-date and not out-of-date. The migration option will be available only if the policy is up-to-date on at least 1 device. Migration does not take place for those devices that have out-of-date policies.
- If the policy is configured in both FlexConfig and the management center:
  - Migration will not be initiated if the policy is already configured at **Device (Edit)** > **Routing**.
  - During deployment, the management center displays an error message. Example EIGRP migration error message—*EIGRP is configured through FlexConfig object and also under Device Listing ->Routing EIGRP for the device. Maintain the EIGRP configuration in either Routing EIGRP or FlexConfig.*
- If the network objects used in the policy exist in the management center, during migration, they are reused. During migration, when a network object that is matching the IP configuration is not available, a new network object is created as *bb* appended with a timestamp and integer, like ***bb_<timestamp>_<integer>***. For more than one such network object, the integer variable in the name would be incremented by one.

Follow these steps to migrate FlexConfig policies:

Procedure

Step 1 Choose **Devices** > **FlexConfig**, click **Edit** (✎) against the FlexConfig policy that you want to migrate.

Step 2 Click **Migrate Config**.

Note
Once migration starts, both **Migrate Config** and FlexConfig **Edit** options are not available.

The **Migrate Config** option is not available, when:

- There are no applicable FlexConfig CLIs to be migrated.
- The FlexConfig policy is not associated with any FlexConfig objects.
- There are no devices that are associated with the FlexConfig policy.

Step 3 In the **Migrate Flex Configuration** dialog box, select the device to which you wish to migrate the configuration, and then click **Ok**.

The migration progress is displayed as task notification. When the migration is complete, click the *View Details* link and download the migration report (PDF format).

Step 4 To view the policy changes, choose **System** (⚙) > **Monitoring** > **Audit** and click on the *Flex Config Migration* message.

Step 5 To view the FlexConfig migration report, choose **System** (⚙) > **Monitoring** > **Audit** and click on the *Flex Config Migration* message. To view the complete migration report, click the **Report** icon.

Step 6 Verify the migrated configuration settings in the corresponding **Device (Edit)** > **Routing**. page.

Step 7 To remove the specific policy config from FlexConfig for the device, in the management center, do the following:

a) Identify the migrated FlexConfig policy for the device.

b) Use the copy option and create duplicate of the FlexConfig policy.
c) Remove the corresponding CLI objects from the duplicated FlexConfig policy.
d) Associate the device to the duplicated FlexConfig policy.

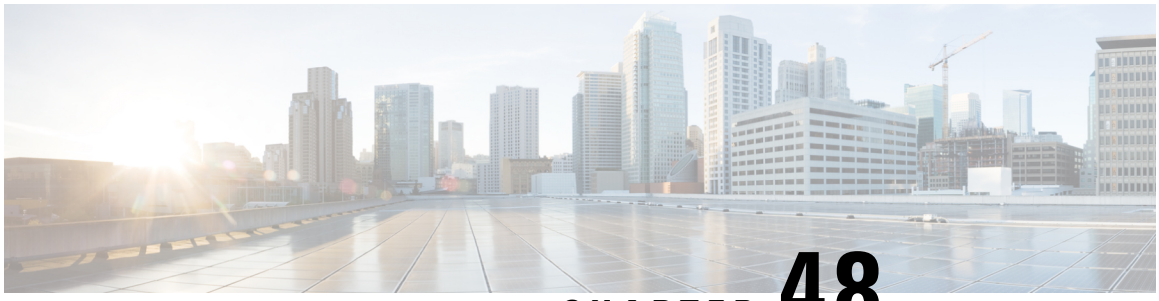
Step 8 Save and deploy the configuration.



PART XIII

Routing

- Static and Default Routes, on page 1357
- Virtual Routers, on page 1375
- ECMP, on page 1431
- Bidirectional Forwarding Detection Routing, on page 1441
- OSPF, on page 1447
- EIGRP, on page 1479
- BGP, on page 1491
- RIP, on page 1513
- Multicast, on page 1521
- Policy-Based Routing, on page 1541



CHAPTER 48

Static and Default Routes

This chapter describes how to configure static and default routes on the Firewall Threat Defense.

- Static and default routes overview, on page 1357
- Requirements and prerequisites for static routes, on page 1360
- Guidelines for static and default routes, on page 1360
- Add a static route, on page 1361
- Reference for routing, on page 1362

Static and default routes overview

To route traffic to a non-connected host or network, define a route to the host or network, using static or dynamic routing. Generally, you must configure at least one static route: a default route for all traffic that is not routed by other means to a default network gateway, which is typically the next-hop router.

Default routes

The simplest option is to configure a default static route to send all traffic to an upstream router, relying on the router to route the traffic for you. A default route identifies the gateway IP address to which the Firewall Threat Defense sends all IP packets for which it does not have a learned or static route. A default static route is simply a static route with 0.0.0.0/0 (IPv4) or ::/0 (IPv6) as the destination IP address.

Default route configuration requirements

You should always define a default route.

The Firewall Threat Defense has separate routing tables for data interfaces and for management-only interfaces (including the special Linux Management interface). You can only add a default route for the data routing table. The Firewall Threat Defense automatically adds a default route in the management-only routing table that sends traffic to the Linux Management interface, where a separate route lookup occurs in the Linux routing table. You can add static routes to the Linux routing table that can be used by Management using the Firewall Threat Defense CLI **configure network static-routes** command.

Note The *default* Linux route is set with the **configure network ipv4** or **configure network ipv6** command.

Static routes

You might want to use static routes in the following cases:

Use cases for static routes

- Your networks use an unsupported router discovery protocol.
- Your network is small and you can easily manage static routes.
- You do not want the traffic or CPU overhead associated with routing protocols.
- In some cases, a default route is not enough. The default gateway might not be able to reach the destination network, so you must also configure more specific static routes. For example, if the default gateway is outside, then the default route cannot direct traffic to any inside networks that are not directly connected to the Firewall Threat Defense.
- You are using a feature that does not support dynamic routing protocols.
- Virtual routers use static routes to create route leaks. Route leaks enable flow of traffic from an interface of a virtual router to another interface in another virtual router. For more information, see Virtual router interconnection, on page 1378.

Route to null0 interface to drop unwanted traffic

Access rules let you filter packets based on the information contained in their headers. A static route to the null0 interface is a complementary solution to access rules. You can use a null0 route to forward unwanted or undesirable traffic so the traffic is dropped.

Performance and routing benefits

Using static null0 routes provides a favorable performance profile. You can also use static null0 routes to prevent routing loops. BGP can leverage the static null0 route for Remotely Triggered Black Hole routing.



Note Though routing traffic to a null0 interface to blackhole it helps to drop unwanted traffic, this method increases the `no-route` ASP drops.

Route priorities

- Routes that identify a specific destination take precedence over the default route.
- When multiple routes exist to the same destination (either static or dynamic), then the administrative distance for the route determines priority. Static routes are set to 1, so they typically are the highest priority routes.
- When you have multiple static routes to the same destination with the same administrative distance, see Equal-cost multi-path routing, on page 1371.
- For traffic emerging from a tunnel with the Tunneled option, this route overrides any other configured or learned default routes.

Transparent firewall mode and bridge group routes

For traffic that originates on the Firewall Threat Defense device and is destined through a bridge group member interface for a non-directly connected network, you need to configure either a default route or static routes so the Firewall Threat Defense device knows out of which bridge group member interface to send traffic. Traffic that originates on the Firewall Threat Defense device might include communications to a syslog server or SNMP server. If you have servers that cannot all be reached through a single default route, then you must configure static routes. For transparent mode, you cannot specify the BVI as the gateway interface; only member interfaces can be used. For bridge groups in routed mode, you must specify the BVI in a static route; you cannot specify a member interface. Refer to MAC address and route lookups, on page 159 for more information.

Static route tracking

The static route tracking feature provides a method for tracking the availability of a static route and installing a backup route if the primary route should fail. For example, you can define a default route to an ISP gateway and configure a backup default route to a secondary ISP in case the primary ISP becomes unavailable.

Static route tracking implementation

One of the problems with static routes is that there is no inherent mechanism for determining if the route is up or down. They remain in the routing table even if the next hop gateway becomes unavailable. Static routes are only removed from the routing table if the associated interface on the Firewall Threat Defense device goes down.

The Firewall Threat Defense device implements static route tracking by associating a static route with a monitoring target host on the destination network that the Firewall Threat Defense device monitors using ICMP echo requests. If an echo reply is not received within a specified time period, the host is considered down, and the associated route is removed from the routing table. An untracked backup route with a higher metric is used in place of the removed route.

When selecting a monitoring target, you need to make sure that it can respond to ICMP echo requests. You can choose any network object as the target, but consider using one of these options:

- The address of the ISP gateway (for dual ISP support)
- The address of the next hop gateway (if you are concerned about gateway availability)
- A server on the target network, such as a syslog server, with which the Firewall Threat Defense device needs to communicate
- A persistent network object on the destination network

Note A PC that may be shut down at night is not a good choice.

You can configure static route tracking for statically defined routes or default routes obtained through DHCP or PPPoE. You can only enable PPPoE clients on multiple interfaces with route tracking configured.

Requirements and prerequisites for static routes

Requirement	Description
Model support	Firewall Threat Defense
Supported domains	Any
User roles	Admin Network Admin

Guidelines for static and default routes

Firewall mode and bridge groups

Configure static routes according to your firewall mode and interface requirements.

- In transparent mode, static routes must use the bridge group member interface as the gateway; you cannot specify the BVI.
- In routed mode, you must specify the BVI as the gateway; you cannot specify the member interface.
- Static route tracking is not supported for bridge group member interfaces or on the BVI.

Supported network address

Ensure your network addresses meet static route requirements.

- Static route tracking is not supported for IPv6.
- The Firewall Threat Defense does not support Class E routing, so a Class E network is not routable in static routes.

Clustering

Consider these limitations for clustering and multiple context mode configurations.

- In clustering, static route tracking is only supported on the control node.

Network object group

You cannot use a range of network objects or a network object group having a range of IP addresses in a static route.

ASP and RIB Route Entries

All routes and its distance installed on the device are captured in the ASP routing table. This is common for all static and dynamic routing protocols. Only the best distance route is captured in the RIB table.

Add a static route

A static route defines where to send traffic for specific destination networks. You should at a minimum define a default route. A default route is simply a static route with 0.0.0.0/0 as the destination IP address.

To configure routes for redundant manager access data interfaces, see Configure a redundant manager access data interface, on page 1153.

Procedure

Step 1 To configure routes for redundant manager access data interfaces, see Configure a redundant manager access data interface, on page 1153.

Step 2 Click **Routing**.

Step 3 (May be required) From the virtual routers drop-down list, select the virtual router for which you are configuring a static route.

Step 4 Select **Static Route**.

Step 5 Click **Add Routes**.

Step 6 Click **IPv4** or **IPv6** depending on the type of static route that you are adding.

Step 7 Choose the **Interface** to which this static route applies.

For transparent mode, choose a bridge group member interface name. For routed mode with bridge groups, you can choose either the bridge group member interface for the BVI name. To "black hole" unwanted traffic, choose the **Null0** interface.

For a device using virtual routing, you can select an interface that belongs to another virtual router. You can create such a static route if you want to leak traffic from this virtual router into the other virtual router. For more information, see Virtual router interconnection, on page 1378.

Step 8 In the **Available Network** list, choose the destination network.

To define a default route, create an object with the address 0.0.0.0/0 and select it here.

Note
Although you can create and choose a Network Object Group containing a range of IP addresses, the Cloud-Delivered Firewall Management Center does not support using ranges in static routes.

Step 9 In the **Gateway** or **IPv6 Gateway** field, enter or choose the gateway router that is the next hop for this route. You can provide an IP address or a Networks/Hosts object.

When you are using static-route configuration for virtual routers to leak routes, do not specify the next hop gateway.

Step 10 In the **Metric** field, enter the number of hops to the destination network. Valid values range from 1 to 255; the default value is 1.

The metric is a measurement of the "expense" of a route, based on the number of hops (hop count) to the network on which a specific host resides. Hop count is the number of networks that a network packet must traverse, including the destination network, before it reaches its final destination. The metric is used to compare routes among different routing protocols. The default administrative distance for static routes is 1, giving it precedence over routes discovered by dynamic routing protocols but not directly connected routes. The default

administrative distance for routes discovered by OSPF is 110. If a static route has the same administrative distance as a dynamic route, the static route takes precedence. Connected routes always take precedence over static or dynamically discovered routes.

Step 11 (Optional) For a default route, click the **Tunneled** checkbox to define a separate default route for VPN traffic.

You can define a separate default route for VPN traffic if you want your VPN traffic to use a different default route than your non VPN traffic. For example, traffic incoming from VPN connections can be easily directed towards internal networks, while traffic from internal networks can be directed towards the outside. When you create a default route with the tunneled option, all traffic from a tunnel terminating on the device that cannot be routed using learned or static routes, is sent to this route. You can configure only one default tunneled gateway per device. ECMP for tunneled traffic is not supported.

Step 12 (IPv4 static route only) To monitor route availability, enter or choose the name of an SLA (service level agreement) Monitor object that defines the monitoring policy, in the **Route Tracking** field.

See Configure an SLA monitor, on page 3034.

Step 13 Click **Ok**.

Reference for routing

This section describes underlying concepts of how routing behaves within the Firewall Threat Defense.

Path determination

Routing protocols use metrics to evaluate what path will be the best for a packet to travel. A metric is a standard of measurement, such as path bandwidth, that is used by routing algorithms to determine the optimal path to a destination. To aid the process of path determination, routing algorithms initialize and maintain routing tables, which include route information. Route information varies depending on the routing algorithm used.

Path determination process

Routing algorithms fill routing tables with a variety of information. Destination or next hop associations tell a router that a particular destination can be reached optimally by sending the packet to a particular router representing the next hop on the way to the final destination. When a router receives an incoming packet, it checks the destination address and attempts to associate this address with a next hop.

Routing tables also can include other information, such as data about the desirability of a path. Routers compare metrics to determine optimal routes, and these metrics differ depending on the design of the routing algorithm used.

Routers communicate with one another and maintain their routing tables through the transmission of a variety of messages. The routing update message is one such message that generally consists of all or a portion of a routing table. By analyzing routing updates from all other routers, a router can build a detailed picture of network topology. A link-state advertisement, another example of a message sent between routers, informs other routers of the state of the sender links. Link information also can be used to build a complete picture of network topology to enable routers to determine optimal routes to network destinations.

Supported route types

There are several route types that a router can use. The Firewall Threat Defense device uses the following route types:

- Static Versus Dynamic
- Single-Path Versus Multipath
- Flat Versus Hierarchical
- Link-State Versus Distance Vector

Static versus dynamic routing algorithms

Static routing algorithms are actually table mappings established by the network administrator. These mappings do not change unless the network administrator alters them. Algorithms that use static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.

Dynamic routing algorithms

Because static routing systems cannot react to network changes, they generally are considered unsuitable for large, constantly changing networks. Most of the dominant routing algorithms are dynamic routing algorithms, which adjust to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages permeate the network, stimulating routers to rerun their algorithms and change their routing tables accordingly.

Dynamic routing algorithms can be supplemented with static routes where appropriate. A router of last resort (a default route for a router to which all unrouteable packets are sent), for example, can be designated to act as a repository for all unrouteable packets, ensuring that all messages are at least handled in some way.

Single-path versus multipath

Some sophisticated routing protocols support multiple paths to the same destination. Unlike single-path algorithms, these multipath algorithms permit traffic multiplexing over multiple lines. The advantages of multipath algorithms are substantially better throughput and reliability, which is generally called load sharing.

Flat versus hierarchical routing

Some routing algorithms operate in a flat space, while others use routing hierarchies.

In a flat routing system, the routers are peers of all others. In a hierarchical routing system, some routers form what amounts to a routing backbone. Packets from non-backbone routers travel to the backbone routers, where they are sent through the backbone until they reach the general area of the destination. At this point, they travel from the last backbone router through one or more non-backbone routers to the final destination.

Routing systems often designate logical groups of nodes, called domains, autonomous systems, or areas.

In hierarchical systems, some routers in a domain can communicate with routers in other domains, while others can communicate only with routers within their domain. In very large networks, additional hierarchical levels may exist, with routers at the highest hierarchical level forming the routing backbone.

The primary advantage of hierarchical routing is that it mimics the organization of most companies and therefore supports their traffic patterns well. Most network communication occurs within small company groups (domains). Because intradomain routers need to know only about other routers within their domain, their routing algorithms can be simplified, and, depending on the routing algorithm being used, routing update traffic can be reduced accordingly.

Link-state versus distance vector

Link-state algorithms (also known as shortest path first algorithms) flood routing information to all nodes in the internetwork. Each router, however, sends only the portion of the routing table that describes the state of its own links. In link-state algorithms, each router builds a picture of the entire network in its routing tables.

Distance vector algorithms (also known as Bellman-Ford algorithms) call for each router to send all or some portion of its routing table, but only to its neighbors. In essence, link-state algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers. Distance vector algorithms know only about their neighbors. Typically, link-state algorithms are used in conjunction with OSPF routing protocols.

Supported internet protocols for routing

The Firewall Threat Defense device supports several internet protocols for routing. Each protocol is briefly described in this section.

The device supports these internet routing protocols:

- **Enhanced Interior Gateway Routing Protocol (EIGRP)**

  EIGRP is a Cisco proprietary protocol that provides compatibility and seamless interoperation with IGRP routers. An automatic-redistribution mechanism allows IGRP routes to be imported into Enhanced IGRP, and vice versa, so it is possible to add Enhanced IGRP gradually into an existing IGRP network.

- **Open Shortest Path First (OSPF)**

  OSPF is a routing protocol developed for Internet Protocol (IP) networks by the interior gateway protocol (IGP) working group of the Internet Engineering Task Force (IETF). OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area includes an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

- **Routing Information Protocol (RIP)**

  RIP is a distance-vector protocol that uses hop count as its metric. RIP is widely used for routing traffic in the global Internet and is an interior gateway protocol (IGP), which means that it performs routing within a single autonomous system.

- **Border Gateway Protocol (BGP)**

  BGP is an interautonomous system routing protocol. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP). Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

Routing tables

The Firewall Threat Defense uses separate routing tables for data traffic (through-the-device) and for management traffic (from-the-device). This section describes how the routing tables work. For information about the management routing table, see also Routing table for management traffic, on page 1370.

How the routing table is populated

The Firewall Threat Defense routing table can be populated by statically defined routes, directly connected routes, and routes discovered by the dynamic routing protocols. Because the Firewall Threat Defense device can run multiple routing protocols in addition to having static and connected routes in the routing table, it is possible that the same route is discovered or entered in more than one manner.

Route selection criteria

When two routes to the same destination are put into the routing table, the one that remains in the routing table is determined as follows:

- If the two routes have different network prefix lengths (network masks), then both routes are considered unique and are entered into the routing table. The packet forwarding logic then determines which of the two to use.

  For example, if the RIP and OSPF processes discovered the following routes:

  - RIP: 192.168.32.0/24
  - OSPF: 192.168.32.0/19

  Even though OSPF routes have the better administrative distance, both routes are installed in the routing table because each of these routes has a different prefix length (subnet mask). They are considered different destinations and the packet forwarding logic determines which route to use.

- If the Firewall Threat Defense device learns about multiple paths to the same destination from a single routing protocol, such as RIP, the route with the better metric (as determined by the routing protocol) is entered into the routing table.

  Metrics are values associated with specific routes, ranking them from most preferred to least preferred. The parameters used to determine the metrics differ for different routing protocols. The path with the lowest metric is selected as the optimal path and installed in the routing table. If there are multiple paths to the same destination with equal metrics, load balancing is done on these equal cost paths.

- If the Firewall Threat Defense device learns about a destination from more than one routing protocol, the administrative distances of the routes are compared, and the routes with lower administrative distance are entered into the routing table.

Administrative Distances for Routes

You can change the administrative distances for routes discovered by or redistributed into a routing protocol. If two routes from two different routing protocols have the same administrative distance, then the route with the lower *default* administrative distance is entered into the routing table. In the case of EIGRP and OSPF routes, if the EIGRP route and the OSPF route have the same administrative distance, then the EIGRP route is chosen by default.

Administrative distance is a route parameter that the Firewall Threat Defense uses to select the best path when there are two or more different routes to the same destination from two different routing protocols. Because

the routing protocols have metrics based on algorithms that are different from the other protocols, it is not always possible to determine the best path for two routes to the same destination that were generated by different routing protocols.

Each routing protocol is prioritized using an administrative distance value. The following table shows the default administrative distance values for the routing protocols supported by the Firewall Threat Defense.

Table 92: Default Administrative Distance for Supported Routing Protocols

Route Source	Default Administrative Distance
Connected interface	0
VPN route	1
Static route	1
EIGRP Summary Route	5
External BGP	20
Internal EIGRP	90
OSPF	110
IS-IS	115
RIP	120
EIGRP external route	170
Internal and local BGP	200
Unknown	255

The smaller the administrative distance value, the more preference is given to the protocol. For example, if the Firewall Threat Defense receives a route to a certain network from both an OSPF routing process (default administrative distance - 110) and a RIP routing process (default administrative distance - 120), the Firewall Threat Defense chooses the OSPF route because OSPF has a higher preference. In this case, the router adds the OSPF version of the route to the routing table.

A VPN advertised route (V-Route/RRI)) is equivalent to a static route with the default administrative distance 1. But it has a higher preference as with the network mask 255.255.255.255.

In this example, if the source of the OSPF-derived route was lost (for example, due to a power shutdown), the Firewall Threat Defense would then use the RIP-derived route until the OSPF-derived route reappears.

The administrative distance is a local setting. For example, if you change the administrative distance of routes obtained through OSPF, that change would only affect the routing table for the Firewall Threat Defense on which the command was entered. The administrative distance is not advertised in routing updates.

Administrative distance does not affect the routing process. The routing processes only advertise the routes that have been discovered by the routing process or redistributed into the routing process. For example, the RIP routing process advertises RIP routes, even if routes discovered by the OSPF routing process are used in the routing table.

Backup Dynamic and Floating Static Routes

A backup route is registered when the initial attempt to install the route in the routing table fails because another route was installed instead. If the route that was installed in the routing table fails, the routing table maintenance process calls each routing protocol process that has registered a backup route and requests them to reinstall the route in the routing table. If there are multiple protocols with registered backup routes for the failed route, the preferred route is chosen based on administrative distance.

Because of this process, you can create floating static routes that are installed in the routing table when the route discovered by a dynamic routing protocol fails. A floating static route is simply a static route configured with a greater administrative distance than the dynamic routing protocols running on the Firewall Threat Defense. When the corresponding route discovered by a dynamic routing process fails, the static route is installed in the routing table.

How forwarding decisions are made

A forwarding decision is a network routing process that

- determines the appropriate interface for packet transmission based on destination address matching in the routing table
- prioritizes routes with longer network prefix lengths when multiple matches exist, and
- directs packets through default routes when no specific routing table entries match the destination.

Forwarding decision process

Forwarding decisions are made as follows:

- If the destination does not match an entry in the routing table, the packet is forwarded through the interface specified for the default route. If a default route has not been configured, the packet is discarded.
- If the destination matches a single entry in the routing table, the packet is forwarded through the interface associated with that route.
- If the destination matches more than one entry in the routing table, then the packet is forwarded out of the interface associated with the route that has the longer network prefix length.

Note When a route change occurs, the Threat Defense re-evaluates the forwarding decision of established connections the next time a packet is seen on the flow.

The resulting behavior is determined by your configuration:

Configuration	Behavior	Outcome
Default (timeout floating-conn 0, no ECMP zone)	Retains the original egress path initially.	If that path can no longer resolve the destination, the flow is torn down after the conn-holddown period (default 15s) rather than waiting for an idle timeout.
Floating-Connection (timeout floating-conn configured)	Identifies a new optimal path on a different interface (outside the same ECMP zone).	The flow is torn down after the floating-conn timeout to allow for a fresh connection. Note: May trigger %FTD-6-110003 syslog messages.
ECMP Zone (Both interfaces in the same ECMP zone)	Dynamically updates the flow's egress interface.	The flow is switched to the new path atomically and non-disruptively.

Longest prefix match example

For example, a packet destined for 192.168.32.1 arrives on an interface with these routes in the routing table:

- 192.168.32.0/24 gateway 10.1.1.2
- 192.168.32.0/19 gateway 10.1.1.3

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.2, because 192.168.32.1 falls within the 192.168.32.0/24 network. It also falls within the other route in the routing table, but 192.168.32.0/24 has the longest prefix within the routing table (24 bits verses 19 bits). Longer prefixes are always preferred over shorter ones when forwarding a packet.

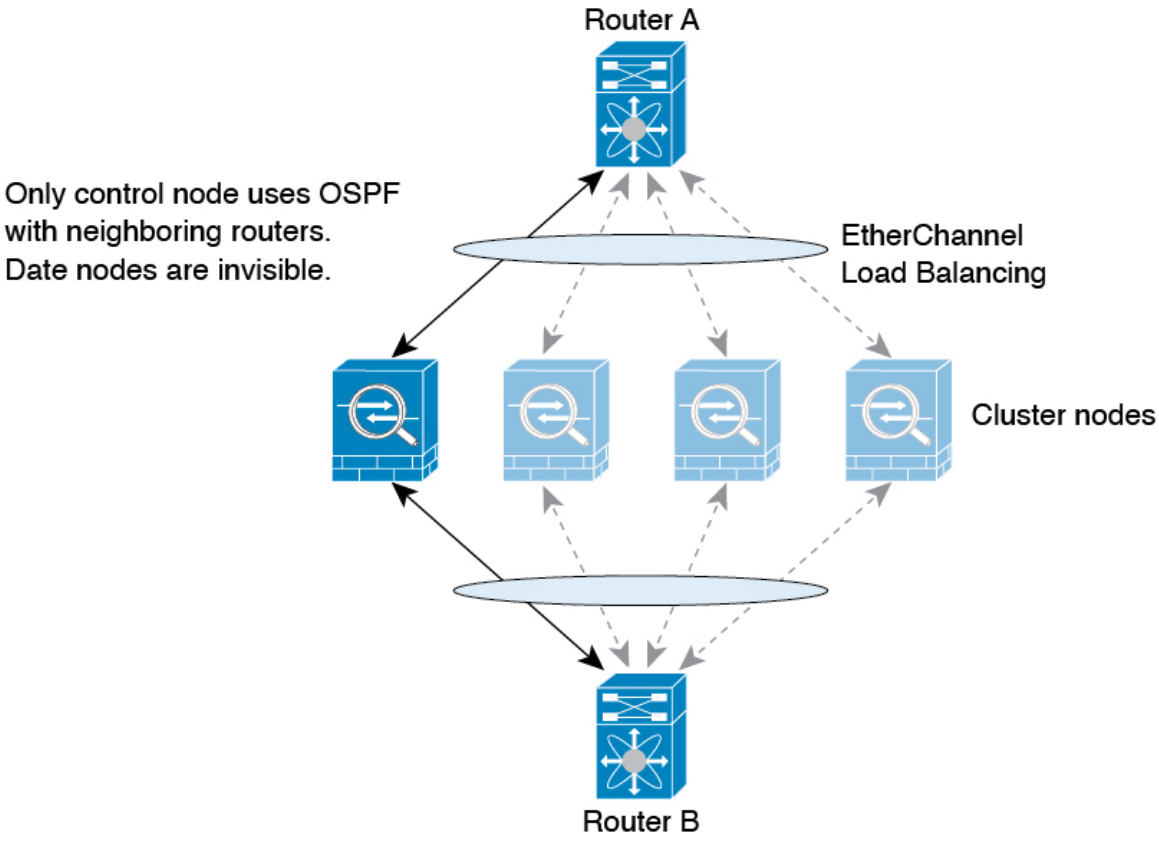
Dynamic routing and High availability

Dynamic routes are synchronized on the standby unit when the routing table changes on the active unit. This means that all additions, deletions, or changes on the active unit are immediately propagated to the standby unit. If the standby unit becomes active in an active/standby ready High availability pair, it will already have an identical routing table as that of the former active unit because routes are synchronized as a part of the High availability bulk synchronization and continuous replication processes.

Dynamic routing in Clustering

The routing process only runs on the control node, and routes are learned through the control node and replicated to data nodes. If a routing packet arrives at a data node, it is redirected to the control node.

Figure 324: Dynamic Routing in Spanned EtherChannel Mode



After the data node learn the routes from the control node, each node makes forwarding decisions independently.

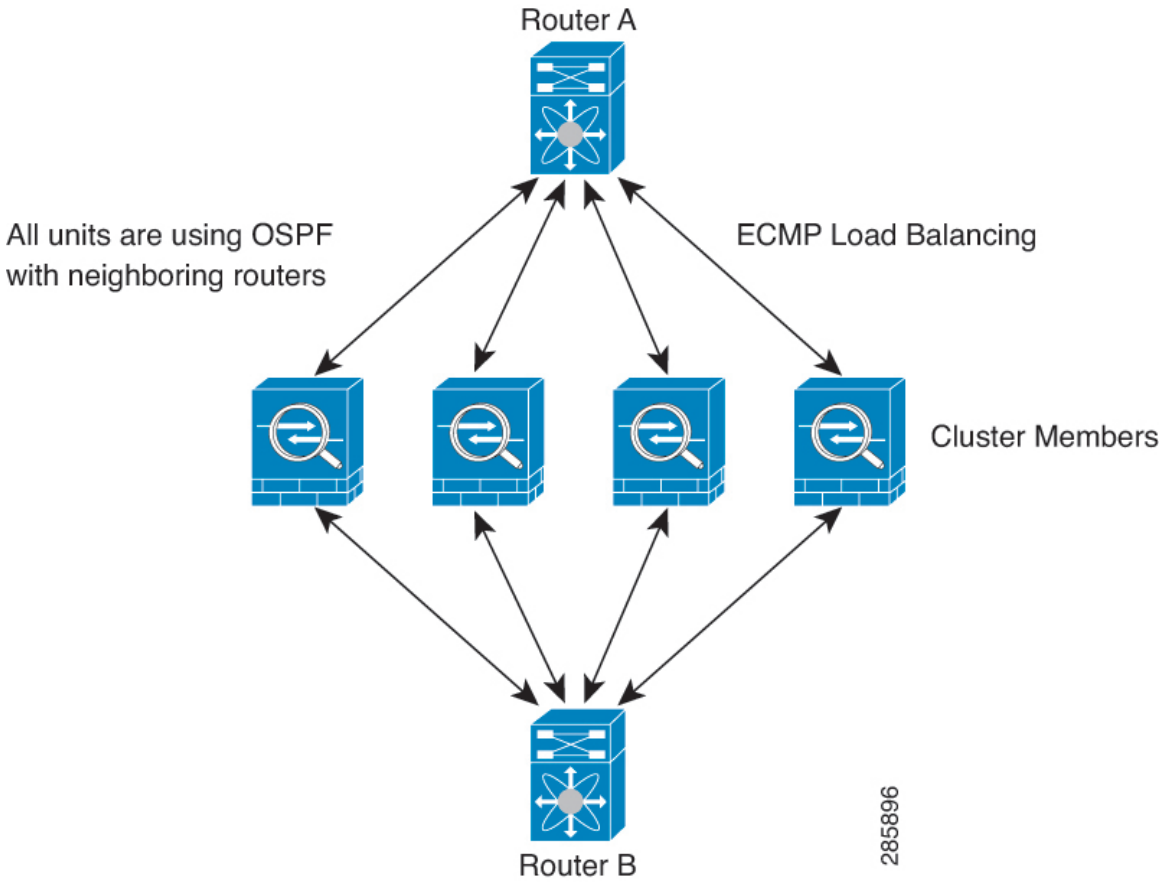
The OSPF LSA database is not synchronized from the control node to data nodes. If there is a control node switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

Dynamic routing in individual interface mode

Learn how dynamic routing functions in individual interface mode, where each cluster node operates as a standalone router. Understand how to configure unique router IDs and manage traffic zones to ensure effective load balancing and prevent asymmetric routing in your network environment.

In Individual interface mode, each node runs the routing protocol as a standalone router, and routes are learned by each node independently.

Figure 325: Dynamic Routing in Individual Interface Mode



In the above diagram, Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when talking to external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

EIGRP does not form neighbor relationships with cluster peers in individual interface mode.

Note If the cluster has multiple adjacencies to the same router for redundancy purposes, asymmetric routing can lead to unacceptable traffic loss. To avoid asymmetric routing, group all of these node interfaces into the same traffic zone. See Create an ECMP zone, on page 1433.

Routing table for management traffic

The Firewall Threat Defense device includes the following routing tables for from-the-device management traffic:

- Linux Management routing table—Special management traffic sourced from the Management interface such as Cloud-Delivered Firewall Management Center communication, licensing communication, and database updates always uses the Linux Management routing table.
- Data routing table—All from-the-device traffic (as well as all through traffic) uses the data routing table by default. All regular data interfaces are part of this routing table. Most services let you choose a specific interface, so only routes associated with that interface are used.
- Management-only routing table—The Management interface and all data interfaces that you set to management-only are part of this routing table. To send from-the-device traffic from any of these interfaces, you must choose a specific management-only interface when you configure the service. An exception is for DNS lookups and ICMP (ping and traceroute): in these cases, the Firewall Threat Defense will use data and then fall back to management automatically if a route is not found. You can add static routes for management-only interfaces, but not for the special Management interface. The Firewall Threat Defense device automatically adds a default route for Management that forwards traffic to Linux, where a separate route lookup occurs in the Linux routing table. You can add static routes to the Linux routing table that can be used by Management using the Firewall Threat Defense CLI **configure network static-routes** command.



Note The *default* Linux route is set with the **configure network ipv4** or **configure network ipv6** command.



Note For devices that have not yet merged the Management and legacy Diagnostic interfaces, see refer to pre-7.3 versions of this guide.

Equal-cost multi-path routing

The Firewall Threat Defense supports Equal-Cost Multi-Path (ECMP) routing.

ECMP routing implementation

You can have up to 8 equal cost static or dynamic routes per interface. For example, you can configure multiple default routes on the outside interface that specify different gateways.

```
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.3
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.4
```

In this case, traffic is load-balanced on the outside interface between 10.1.1.2, 10.1.1.3, and 10.1.1.4. Traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses, incoming interface, protocol, source and destination ports.

ECMP Across Multiple Interfaces Using Traffic Zones

If you configure traffic zones to contain a group of interfaces, you can have up to 8 equal cost static or dynamic routes across up to 8 interfaces within each zone. For example, you can configure multiple default routes across three interfaces in the zone:

```
route for 0.0.0.0 0.0.0.0 through outside1 to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside2 to 10.2.1.2
route for 0.0.0.0 0.0.0.0 through outside3 to 10.3.1.2
```

Similarly, your dynamic routing protocol can automatically configure equal cost routes. The Firewall Threat Defense load-balances traffic across the interfaces with a more robust load balancing mechanism.

When a route is lost, the device seamlessly moves the flow to a different route.

Routing GRE traffic

In Firewall Threat Defense, Generic Routing Encapsulation (GRE) traffic is managed using bi-directional flows. This means that each GRE session is represented by a single flow entry that handles both inbound and outbound traffic between endpoints. With bi-directional flows, only one route lookup is performed and both directions are mapped to the same interface. As a result, the system cannot make independent routing decisions for each direction, which may lead to sub-optimal routing in some scenarios.

About route maps

Route maps are used when redistributing routes into an OSPF, RIP, EIGRP or BGP routing process. They are also used when generating a default route into an OSPF routing process. A route map defines which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process.

Route map characteristics

Route maps have many features in common with widely known ACLs. These are some of the traits common to both:

- They are an ordered sequence of individual statements, and each has a permit or deny result. Evaluation of an ACL or a route map consists of a list scan, in a predetermined order, and an evaluation of the criteria of each statement that matches. A list scan is aborted once the first statement match is found and an action associated with the statement match is performed.
- They are generic mechanisms. Criteria matches and match interpretation are dictated by the way that they are applied and the feature that uses them. The same route map applied to different features might be interpreted differently.

These are some of the differences between route maps and ACLs:

- Route maps are more flexible than ACLs and can verify routes based on criteria which ACLs can not verify. For example, a route map can verify if the type of route is internal.
- Each ACL ends with an implicit deny statement, by design convention. If the end of a route map is reached during matching attempts, the result depends on the specific application of the route map. Route maps that are applied to *redistribution* behave the same way as ACLs: if the route does not match any clause in a route map then the route redistribution is denied, as if the route map contained a deny statement at the end.

Permit and deny clauses

Route maps can have permit and deny clauses. The deny clause rejects route matches from redistribution. You can use an ACL as the matching criterion in the route map.

Route map and ACL clause interaction rules

Because ACLs also have permit and deny clauses, the following rules apply when a packet matches the ACL:

- ACL permit + route map permit: routes are redistributed.
- ACL permit + route map deny: routes are not redistributed.
- ACL deny + route map permit or deny: the route map clause is not matched, and the next route-map clause is evaluated.

Match and set clause values

Each route map clause has two types of values:

- A match value selects routes to which this clause should be applied.
- A set value modifies information that will be redistributed into the target protocol.

Route map evaluation process

For each route that is being redistributed, the router first evaluates the match criteria of a clause in the route map. If the match criteria succeeds, then the route is redistributed or rejected as dictated by the permit or deny clause, and some of its attributes might be modified by the values set from the set commands. If the match criteria fail, then this clause is not applicable to the route, and the software proceeds to evaluate the route against the next clause in the route map. Scanning of the route map continues until a clause is found that matches the route or until the end of the route map is reached.

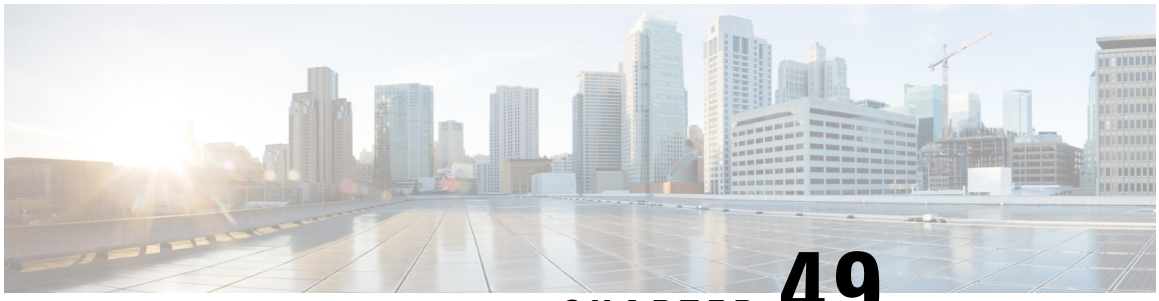
A match or set value in each clause can be missed or repeated several times, if one of these conditions exists:

- If several match entries are present in a clause, all must succeed for a given route in order for that route to match the clause (in other words, the logical AND algorithm is applied for multiple match commands).
- If a match entry refers to several objects in one entry, either of them should match (the logical OR algorithm is applied).
- If a match entry is not present, all routes match the clause.
- If a set entry is not present in a route map permit clause, then the route is redistributed without modification of its current attributes.



Note Do not configure a set entry in a route map deny clause because the deny clause prohibits route redistribution—there is no information to modify.

A route map clause without a match or set entry does perform an action. An empty permit clause allows a redistribution of the remaining routes without modification. An empty deny clause does not allow a redistribution of other routes (this is the default action if a route map is completely scanned, but no explicit match is found).



CHAPTER 49

Virtual Routers

This chapter describes underlying concepts about virtual routers and on how virtual routing behaves within the Secure Firewall Threat Defense.

- Virtual routers and virtual routing and forwarding (VRF), on page 1375
- Maximum number of virtual routers by device model, on page 1381
- Requirements and prerequisites for virtual routers, on page 1383
- Guidelines and limitations for virtual routers, on page 1383
- Manage virtual routers, on page 1386
- Create a virtual router, on page 1386
- Monitor virtual routers, on page 1390
- Configuration examples for virtual routers, on page 1390

Virtual routers and virtual routing and forwarding (VRF)

A virtual router is a network virtualization feature that

- maintains separate routing tables for groups of interfaces
- provides clean separation in the traffic flowing through the device, and
- implements the "light" version of Virtual Routing and Forwarding, or VRF-Lite, which does not support Multiprotocol Extensions for BGP (MBGP).

Virtual router capabilities and configuration

Virtual routers provide support to two or more distinct customers over a common set of networking equipment. You can also use virtual routers to provide more separation for elements of your own network, for example, by isolating a development network from your general purpose corporate network.

When you create a virtual router, you assign interfaces to the router. You can assign a given interface to one, and only one, virtual router. You would then define static routes, and configure routing protocols such as OSPF or BGP, for each virtual router. You would also configure separate routing processes over your entire network, so that routing tables on all participating devices are using the same per-virtual-router routing process and tables. Using virtual routers, you create logically-separated networks over the same physical network to ensure the privacy of the traffic that runs through each virtual router.

Because the routing tables are separate, you can use the same, or overlapping, address spaces across the virtual routers. For example, you could use the 192.168.1.0/24 address space for two separate virtual routers, supported by two separate physical interfaces.

Note that there are separate management and data routing tables per virtual router. For example, if you assign a management-only interface to a virtual router, then the routing table for that interface is separate from the data interfaces assigned to the virtual router.

Virtual routers and dynamic VTI

You can create virtual routers, associate dynamic VTIs with these virtual routers, and extend the capabilities of dynamic VTIs in your network. You can associate dynamic VTIs either with global or user-defined virtual routers. You can assign a dynamic VTI to only one virtual router. For more information, refer to Virtual Routers and Dynamic VTI, on page 1749.

Virtual router applications

Virtual router applications are network implementation scenarios that

- isolate networks on shared resources through dedicated routing tables
- enable common security policy management across different departments or networks, and
- provide shared internet access for different departments or networks.

Global and user-defined virtual routers

This reference provides information about global and user-defined virtual routers, including their characteristics and the features supported by each type.

Global virtual routers

For a device with virtual routing capability, system creates a global virtual router by default. The system assigns all interfaces in your network to the global virtual router. A routed interface can belong to either a user-defined virtual router or a global virtual router. When you upgrade Firewall Threat Defense to a version which has virtual router capability, all its existing routing configuration becomes part of the global virtual router.

User-Defined virtual routers

A user-defined virtual router is the one defined by you. You can create more than one virtual router on a device. However, anytime, an interface can be assigned to only one user-defined virtual router. While some of the device features are supported on user-defined virtual routers, few of the features are supported only on the global virtual routers. User-defined virtual routers support route-based site-to-site VPN (static VTI) (static and dynamic VTI).

Supported features and monitoring policies

You can configure these features on the global virtual router only:

- OSPFv3
- RIP

- EIGRP
- IS-IS
- Multicast Routing
- Policy Based Routing (PBR)

ISIS and PBR are supported through Flex Config in Cloud-Delivered Firewall Management Center (see Predefined FlexConfig objects, on page 1322). Configure only global virtual router interfaces for these features.

DHCP server auto-configuration uses WINS/DNS server that is learned from an interface. This interface can only be a global virtual router interface.

You can configure these features separately for each user-defined virtual router:

- Static routes and their SLA monitors
- OSPFv2
- BGPv4/v6
- Integrated Routing and Bridging (IRB)
- SNMP

These features are used by the system when querying or communicating with the remote system (from-the-box traffic). These features use interfaces in the global virtual router only. That means, if you configure an interface for the feature, it must belong to the global virtual router. As a rule, anytime, if the system must look up a route to reach an external server for its own management purposes, it does the route lookup in the global virtual router.

- DNS server when used to resolve fully qualified names used in access control rules, or for resolving names for the **ping** command. If you specify **any** as the interface for a DNS server, the system considers interfaces in the global virtual router only.
- AAA server or identity realm when used with VPN. You can configure VPN on interfaces in the global virtual router only. Thus, the external AAA servers that are used for VPN, such as Active Directory, must be reachable through an interface in the global virtual router.
-

Virtual-router-aware policies

Virtual-router-aware policies are security policies that

- apply to interfaces from a single virtual router only
- use security zones whose memberships are constrained to interfaces assigned to a specific virtual router, and
- prevent policies from affecting all interfaces across all virtual routers when that is not desired.

Virtual router policy behavior

When you create a virtual router, the routing table for that virtual router is automatically separated from the global virtual router or any other virtual router. However, security policies are not automatically virtual-router-aware.

For example, if you write an access control rule that applies to "any" source or destination security zone, then the rule will apply to all interfaces across all virtual routers. This might in fact be exactly what you want. For example, all of your customers might want to block access to the same list of objectionable URL categories.

But, if you need to apply a policy to one of the virtual routers but not others, you need to create security zones that contain interfaces from that single virtual router only. Then, use the virtual-router-constrained security zones in the source and destination criteria of the security policy.

By using security zones whose memberships are constrained to the interfaces assigned to a single virtual router, you can write virtual-router-aware rules in these policies:

- Access control policy.
- Intrusion and file policies.
- SSL decryption policy.
- Identity policy and user-to-IP address mappings. If you use overlapping address spaces in virtual routers, ensure that you create separate realms for each virtual router and apply them correctly in the identity policy rules.

If you use overlapping address spaces in your virtual routers, you should use security zones to ensure that the right policies get applied. For example, if you use the 192.168.1.0/24 address space in two separate virtual routers, an access control rule that simply specifies the 192.168.1.0/24 network will apply to traffic in both virtual routers. If that is not the desired outcome, you can limit the application of the rule by also specifying the source/destination security zones for just one of the virtual routers.

Virtual router interconnection

Virtual router interconnection is a routing mechanism that

- enables traffic routing between virtual routers through route leaking,
- supports both manual static routes and dynamic BGP configurations, and
- allows network segmentation while maintaining controlled inter-segment communication.

Route leaking methods

You can configure the device to route traffic between virtual routers. This process of route leaking can be done manually by setting up static routes or dynamically through BGP settings.

Static route leaking involves manual configuration:

You can configure static routes to route traffic between virtual routers.

For example, if you have the outside interface in the global virtual router, you can set up static default routes in each of the other virtual routers to send traffic to the outside interface. Then, any traffic that cannot be routed within a given virtual router gets sent to the global router for subsequent routing.

Static routes between virtual routers are known as route leaks, because you are leaking traffic to a different virtual router. When you are leaking routes, say, VR1 routes to VR2, you can initiate connections from VR2 to VR1 only. For traffic to flow from VR1 to VR2, you must configure the reverse route. When you create a static route to an interface in another virtual router, you do not need to specify a gateway address. Simply select the destination interface.

For inter-virtual-router routes, the system does destination interface look-up in the source virtual router. Then, it looks up the MAC address of the next hop in the destination virtual router. Thus, the destination virtual router must have either a dynamic (learned) or static route for the selected interface for the destination address.

Configuring NAT rules that use source and destination interfaces in different virtual routers can also allow traffic to route between virtual routers. If you do not select the option for NAT to do a route lookup, the rule will simply send traffic out the destination interface with a NATed address whenever destination translation happens. However, the destination virtual router should have a route for the translated destination IP address so that next-hop lookup can succeed.

Though NAT rule leaks traffic from one virtual router to another, to ensure correct routing, we recommend that you configure a static route leak between these virtual routers for the translated traffic. Without the route leak, sometimes the rule may not match the traffic you expect it to match, and the translation may not be applied.

Virtual routing does not support a cascading or chain of route leaks. For example, assume that your Firewall Threat Defense has VR1, VR2, and VR3 virtual routers; VR3 is directly connected to a network - 10.1.1.0/24. Now, assume you configure a route leak in VR1 for network 10.1.1.0/24 through interface in VR2 and in VR2 define a route leak for 10.1.1.0/24 through VR3. This chain of route leaks will not allow traffic to hop from VR1 to VR2 and then exit from VR3. In case of route leaks, the route lookups first determine egress interface from input Virtual Router's routing table and then looks at the output of Virtual Router's routing table for next hop lookup. From both the lookups, egress interface should match. In our example, the egress interfaces will not be the same and hence the traffic does not pass through.

Use static inter VRF route with caution when the destination network is not a direct-connected subnet of the upstream (outgoing) VR. For example, assume two VRs - VR1 and VR2. While VR1 handles the outgoing traffic which gets the default route from its external peer through BGP or any dynamic routing protocol, and VR2 handles the incoming traffic which is configured with static inter VRF default route with VR1 as the next-hop. When VR1 loses the default route from its peer, VR2 will not able to detect that its upstream (outgoing) VR lost the default route and the traffic is still sent toward VR1 which will eventually get dropped without notifications. In this scenario, we recommend that you configure VR2 with dynamic route leak through BGP.

You can implement an inter-virtual-router route leak by exporting routes from a source virtual router (say VR1) to the source BGP table using route target extended community and then importing the same route target extended community from the source BGP table into the destination BGP table, which in turn is used by the destination virtual router (say, VR2). You can use the route maps for filtering the routes. The routes of global virtual router can also be leaked to user-defined virtual routers and vice versa. The BGP inter-virtual-router route leaking supports both ipv4 and ipv6 prefixes.

For details on configuring BGP route leaking, see Configure BGP route import and export, on page 1509.

BGP route leaking guidelines:

- Ensure that all the routes required for recursion are imported and present in the routing table of the ingress virtual router.
- ECMP is supported per virtual router. Hence, do not configure an ECMP across different virtual routers. The overlapping prefixes imported from different virtual routers cannot form an ECMP. That is, when you attempt to import routes with overlapping addresses from two different virtual routers to other virtual

routers (a global virtual router or n user-defined virtual router), only one route (as per BGP best path algorithm, the first route that was advertised) is imported to the respective virtual routing table. For example, if a network 10.10.0.0/24, connected to VR1 is advertised through BGP to a global virtual router first, and later another network with the same address 10.10.0.0/24, connected to VR2 is also advertised through BGP to global virtual router, only the VR1 network route is imported to the global virtual routing table.

- OSPFv3 is not supported on user-defined virtual routers. Hence, do not configure BGPv6 to leak OSPFv3 user-defined virtual routers to global virtual router. However, you can configure BGPv6 to leak OSPFv3 global virtual router routes to user-defined virtual router through redistribution.
- It is recommended to keep VTI interface, protected internal interfaces (loopback interface if supported for VTI) to be part of same virtual router to prevent the need for route leak.

Best practice for overlapping IP addresses

Virtual router creates multiple instances of routing tables that are independent, thereby, the same or overlapping IP addresses can be used without conflicts. Firewall Threat Defense allows the same network to be part of two or more virtual routers. This involves multiple policies to be applied at the interface or at the virtual router level.

Limitations with overlapping IP addresses

When using an overlapping IP address in multiple virtual routers, to ensure proper application of the policy, you have to modify policies or rules for some of the features. Such features require you to use more specific interface either by splitting existing security zone or using new interface group as the need be.

- Network Map—Modify the network discovery policy to exclude some overlapping IP segments to ensure that there is no overlapping IP address being mapped.
- Identity Policy—The identity feed source cannot differentiate among virtual routers; to overcome this limitation, map overlapping address spaces or virtual routers in different realms.
- Access Policy—Apply rules on specific interfaces to ensure that different policies are applied on overlapping IP segments.
- Prefilter Policy—Apply rules on specific interfaces to ensure that different policies are applied on overlapping IP segments.
- QoS/Rate Limit—Apply rules on specific interfaces to ensure that different policies are applied on overlapping IP segments.
- SSL Policy—Apply rules on specific interfaces to ensure that different policies are applied on overlapping IP segments.

Unsupported features with overlapping IP addresses

These features are not supported with overlapping IP addresses:

- ISE SGT-based Rule in AC Policy—The static security group tag (SGT) to IP address mappings downloaded from Cisco Identity Services Engine (ISE) are not virtual-router-aware. Set up separate ISE systems per virtual router if you need to create different SGT mappings per virtual router. This is not necessary if you intend to map the same IP addresses to the same SGT number in each virtual router.
- Overlapping DHCP server pools are not supported across virtual routers.

- Events and Analytics—Many of the Cloud-Delivered Firewall Management Center analytics are dependent on network map and identity mappings which cannot differentiate if the same IP address belongs to two different end hosts. Hence, these analytics are not accurate when there are overlapping IP segments existing in same device but different virtual routers.

Other than few exceptions, the routing functions and most of the NGFW and IPS capability does not get impacted by the overlapping IP addresses.

Configuring SNMP on user-defined virtual routers

In addition to supporting SNMP on the management interface and Global virtual router data interfaces, Secure Firewall Threat Defense now allows you to configure SNMP host on user-defined virtual routers.

Summary

The key components involved in configuring SNMP on user-defined virtual routers are:

- Device interfaces: Physical or logical interfaces that need to be configured and enabled
- Virtual router: User-defined virtual router configured on specific interfaces
- SNMP hosts: Network management stations that receive SNMP polling and traps
- Secure Firewall Threat Defense: The target device where configurations are deployed

Workflow

The SNMP configuration on user-defined virtual routers involves these stages:

1. Configure device interfaces.
2. Configure virtual router on an interface.
3. Configure SNMP hosts on a virtual router interface.



Note SNMP is not virtual router-aware. Hence, while configuring SNMP server on the user-defined virtual router, ensure that the network address is not an overlapping IP address.

4. Deploy configurations to Secure Firewall Threat Defense.

   On successful deployment, the SNMP polling and traps are sent to the Network Management Station through the virtual router interface.

Maximum number of virtual routers by device model

The maximum number of virtual routers you can create depends on the device model. The following table shows the maximum supported virtual routers per model. You can confirm the maximum for your system by entering the **show vrf counters** command, which displays the maximum number of user-defined virtual routers for your platform (not including the global virtual router). The numbers below include both user and global routers. For Firepower 4100/9300 devices, these values apply to native mode.

For platforms that support multi-instance capability, such as the Firepower 4100/9300, determine the maximum number of virtual routers per container instance by dividing the maximum virtual routers by the number of cores on the device, and then multiplying by the number of cores assigned to the instance, rounding down to the nearest whole number. For example, if the platform supports a maximum of 100 virtual routers, and it has 70 cores, then each core would support a maximum of 1.43 virtual routers (rounded). Thus, an instance assigned 6 cores would support 8.58 virtual routers, which rounds down to 8, and an instance assigned 10 cores would support 14.3 virtual routers (rounding down, 14).

Device Model	Maximum Virtual Routers
Secure Firewall 220	5
Firepower 1010	5
Firepower 1120	5
Firepower 1140	10
Firepower 1150	10
Secure Firewall 1210CE	5
Secure Firewall 1210CP	5
Secure Firewall 1220CX	10
Secure Firewall 1230	10
Secure Firewall 1240	10
Secure Firewall 1250	15
Firepower 2110	10
Firepower 2120	20
Firepower 2130	30
Firepower 2140	40
Secure Firewall 3105	10
Secure Firewall 3110	15
Secure Firewall 3120	25
Secure Firewall 3130	50
Secure Firewall 3140	100
Firepower 4112	60
Firepower 4115	80
Firepower 4125	100

Device Model	Maximum Virtual Routers
Firepower 4145	100
Secure Firewall 4215	100
Secure Firewall 4225	100
Secure Firewall 4245	100
Firepower 9300 appliance, all models	100
Firewall Threat Defense Virtual, all platforms	30
ISA 3000	10

Related Topics

Requirements and prerequisites for container instances

Requirements and prerequisites for virtual routers

This reference provides the essential requirements and prerequisites that must be met before deploying and configuring virtual routers in your network environment.

Model support

Firewall Threat Defense

Supported domains

Any

User roles

Admin

Network Admin

Security Approver

Guidelines and limitations for virtual routers

Firewall mode guidelines

Use virtual routers on routed firewall mode only.

Interface guidelines

Configure interfaces for virtual routers according to these guidelines:

- Assign an interface to only one virtual router.
- A virtual router can have any number of interfaces that are assigned to it.
- Assign only routed interfaces with logical names and VTIs to a user-defined virtual router.
- If you want to change a virtual router interface to a non-routed mode, remove the interface from the virtual router, and then change its mode.
- You can assign an interface to a virtual router, either from a global virtual router or from another user-defined virtual router.
- Do not assign these interfaces to a user-defined virtual router:
  - Members of EtherChannel.
  - Members of Redundant interfaces.
  - Members of BVI.
- VTI is a route-based VPN. When the tunnel is established, control the traffic that uses VTI for encryption through routing. Static routing, as well as dynamic routing with BGP, OSPFv2/v3, or EIGRP is supported.
- Do not use interfaces that belong to user-defined virtual routers in policy-based site-to-site or remote access VPNs.
- A dynamic VTI and its corresponding protected network interface must be part of the same virtual router.
- You must map the borrow IP interface and the dynamic VTI to the same virtual router.
- User-defined virtual routers support only BGPv4/v6 and OSPFv2 routing protocols.
- A tunnel source interface can be in a different user-defined virtual router than that associated with the dynamic VTI.

Follow these precautions when moving interfaces between virtual routers:

- If a route using the interface that is being moved or its virtual router is deleted, exist in source or destination virtual router table, remove the routes before the interface movement or virtual router deletion.
- As separate routing tables are maintained for each virtual router, when an interface is moved from one virtual router to another virtual router, be it global or user-defined, the system removes the IP address configured on the interface temporarily. All existing connections on the interface are terminated. Thus, moving interfaces between virtual routers have drastic effect on the network traffic. Hence take precautionary measures before you move interfaces.

Global virtual router guidelines

Consider these guidelines for global virtual routers:

- The interfaces which are named and not part of other virtual routers, are part of the global virtual router.
- Do not remove routed interfaces from global virtual router.
- Do not modify global virtual router.
- Generally, after configuring interfaces, if you un-register and register back to same or another Cloud-Delivered Firewall Management Center, interface configuration is imported back from device.

With virtual router support, there is a restriction—the IP address for only global virtual router interfaces is retained.

Clustering guidelines

Consider these limitations when using virtual routers with clustering:

- When the control unit link fails due to failure of its interfaces, the unit removes all leaked routes of its interfaces from the global routing table and propagates the inactive connected and static routes to other units of the cluster. This results in removal of those leaked routes from the routing table in other units. These removal takes place prior to another unit becoming a new control unit, which takes approximately 500 ms. When another unit becomes the new control unit, these routes are learned and added back to the routing tables through BGP convergence. Thus, till the convergence time, approximately one minute, the leaked routes are not available for the routing events to take place.
- When a control role change occurs in a cluster, the leaked routes learnt through BGP is updated with the best ECMP path. However, the non-best ECMP path is removed from the cluster routing table only after the BGP reconvergence timer elapses, that is 210 seconds. Thus, till the BGP reconvergence timer elapse, the old, non-best ECMP path persist as preferred route for routing events.

Additional guidelines

Follow these additional guidelines when configuring virtual routers:

- While configuring BGP for virtual routers, you can redistribute the routes belonging to different protocols within the same virtual routers. For example, OSPF VR2 routes cannot be imported into BGP VR1. You can only redistribute OSPF VR2 into BGP VR2, and then configure a route leak between BGP VR2 and BGP VR1.
- Do not use IPv6 ACL for filtering the routes in the route map. Only prefix list is supported.
- Security Intelligence Policy—The Security Intelligence policy is not virtual-router-aware. If you add an IP address, URL, or DNS name to the block list, it is blocked for all virtual routers. This limitation is applicable on the interface having security zones.
- NAT Rules—Do not mix interfaces in NAT rules. In virtual routing, if the specified source and destination interface objects (interface groups or security zones) have interfaces that belong to different virtual routers, the NAT rule diverts traffic from one virtual router through another virtual router. The NAT does the route lookup in the virtual router table for the inbound interface only. If necessary, define static routes in the source virtual router for the destination interface. If you leave the interface as **any**, the rule applies to all interfaces regardless of virtual router membership.
- DHCP Relay—Interconnecting virtual routers for DHCP relay is not supported. For example, if DHCP relay client is enabled on VR1 interface and DHCP relay server is enabled on VR2 interface, the DHCP requests will not be forwarded outside of VR2 interface.
- Recreating a deleted virtual router—When you recreate a virtual router that was deleted less than 10 seconds earlier, an error message appears stating that deletion of the virtual router is in progress. If you want to recreate a deleted virtual router successively, use a different name for the new virtual router.

Manage virtual routers

When you click **Manage Virtual Routers** on the Virtual Routers pane, the Manage Virtual Routers page appears. This page displays the existing virtual routers on the device and the associated interfaces. On this page, you can **Add Virtual Router** (+) to the device. You can also **Edit** (✎) and **Delete** (🗑) user-defined virtual routers. You cannot edit or remove a global virtual router. You can only **View** (🔍) the details of a global virtual router.

Create a virtual router

Create a virtual router to enable network segmentation and manage routing configurations independently within the firewall device.

Virtual routers allow you to segment network traffic and create separate routing domains within a single firewall device, providing enhanced network management and security isolation.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Manage Virtual Routers**.

Step 4 Click **Add Virtual Router** (+).

Step 5 In the Add Virtual Router box, enter a name and description for the virtual router.

Note
If you are creating a virtual router that was deleted less than 10 seconds earlier, an error message appears stating that deletion of the virtual router is in progress. If you want to create a deleted virtual router successively, use a different name for the new virtual router.

Step 6 Click **Ok**.

The Routing page appears, displaying the newly created virtual router page.

The virtual router is successfully created and appears in the routing configuration interface, ready for further configuration.

What to do next

- Configure a Virtual Router.

Configure a virtual router

Configure a virtual router to assign interfaces and set up routing policies for network traffic management and routing capabilities on your device.

You can assign interfaces to a user-defined virtual router and configure the routing policies for the device. Though you cannot manually add or remove interfaces for a global virtual router, you can configure the routing policies for the device interfaces.

Before you begin

- To configure routing policies for a user-defined virtual router, add a router. See Create a virtual router, on page 1386.
- All routing configuration settings of a non-virtual routing capable device are also available for a global virtual router. For information on the settings, see Routing Settings.
- Only limited routing protocols are supported for a user-defined virtual router.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, from the **Devices** > **Device Management** page, edit the virtual-router supported device. Navigate to **Routing**.

For information on the modifications to the routing page, see Modifications to the Cloud-Delivered Firewall Management Center web interface routing page.

Step 2 From the drop-down list, select the desired virtual router.

Step 3 In the **Virtual Router Properties** page, modify the description as needed.

Step 4 To add interfaces, select the interface under the **Available Interfaces** box, and then click **Add**.

Remember these guidelines:

- Only interfaces with a logical name are listed under the **Available Interfaces** box. You can edit the interface and provide a logical name in **Interfaces**. Remember to save the changes for the settings to take effect.
- Only interfaces of global virtual routers are available for assigning; the **Available Interfaces** box lists only interfaces that are not assigned to any other user-defined virtual routers. You can assign physical interfaces, subinterfaces, redundant interfaces, bridge groups, VTIs, and EtherChannels to a virtual router, but not their member interfaces. Because the member interfaces cannot be named, they cannot be used in virtual routing.

  You can assign the diagnostic interface to the global virtual router only.

Step 5 To save the interface settings, click **Save**.

Step 6 To configure the routing policy for the virtual router, click the respective names to open the corresponding settings page:

- **OSPF**—Only OSPFv2 is supported on the user-defined virtual router. All other settings for OSPFv2 are as applicable as for a non-virtual-router-aware interface, except that **Interface** allows you to select only the interfaces of the virtual router that you are configuring. You can define the OSPFv3 and OSPFv2 routing policies for a global virtual router. For information on the OSPF settings, see OSPF, on page 1447.

- **RIP**—You can configure RIP routing policies only for a global virtual router. For information on RIP settings, see RIP, on page 1513.
- **BGP**—This page displays the BGP general settings that you have configured in **Settings**:
  - You cannot modify any of those general settings on this page, except for the router ID settings. You can override the router ID settings that were defined in the **Settings** page by editing them on this page.
  - To configure other BGP IPv4 or IPv6 settings, you must enable the BGP option in **BGP** page under **General Settings**.
  - BGP configuration for both IPv4 and IPv6 address family are supported for global router and the user-defined virtual router.

  For information on configuring BGP settings, see BGP, on page 1491.
- **Static Route**—Use this setting to define where to send traffic for a specific destination network. You can also use this setting to create an inter-virtual-router static route. You can create a leak of connected or static route by using the interfaces of user-defined or global virtual routers. **FMC prefixes** to an interface to indicate that it is belonging to another virtual router and can be used for a route leak. For the route leak to be successful, do not specify next hop gateway.

  The Static Route table displays the virtual router whose interface is used for a route leak in the **Leaked from Virtual Router** column. If it is not a route leak, the column displays N/A.

  Irrespective of which virtual router the static route belongs, a Null0 interface is listed along with the interfaces of the same virtual router to which the static route belongs.

  For information on static route settings, see Static and Default Routes, on page 1357.
- **Multicast**—You can configure multicast routing policies only for a global virtual router. For information on multicast settings, see Multicast, on page 1521.

Step 7 To save the routing policy settings, click **Save**.

The virtual router is configured with the assigned interfaces and routing policies you specified.

What to do next

- Modify a virtual router.
- Remove a virtual router.

Modify a virtual router

This task allows you to update the configuration settings of virtual routers to meet changing network requirements.

You can modify the description and other routing policies of a virtual router.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Manage Virtual Routers**.

All virtual routers along with the assigned interfaces are displayed in the **Virtual Routers** page.

Step 4 To modify a virtual router, click **Edit** ( ) against the desired virtual router.

Note
You cannot modify the general settings of the global virtual router. Hence, edit is not available for the global router; instead a **View** ( ) is provided to view the settings.

Step 5 To save the changes, click **Save**.

The virtual router configuration is updated with your changes.

What to do next

- Remove a virtual router.

Remove Virtual Routers

Remove virtual routers when they are no longer needed or when consolidating routing configurations. This helps maintain a clean and efficient network routing structure.

Virtual routers provide network segmentation and routing isolation. When removing virtual routers, all associated interfaces automatically move to the global virtual router, and all routing policies are deleted.

Before you begin

- You cannot delete the Global virtual router. Hence, the delete option is not available for the Global virtual router.
- You can remove multiple virtual routers at a time.
- All the routing policies of the deleted virtual router are also deleted.
- All the interfaces of the deleted virtual router move to the global virtual router.
- If there are any restrictions on the movement of interfaces, such as overlapping IPs, route conflicts, and so on, you can remove the router only after resolving the conflicts.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Manage Virtual Routers**.

All virtual routers along with the mapped interfaces are displayed in the **Virtual Routers** page.

Step 4 To remove a virtual router, click **Delete** ( ) against the desired virtual router.

Step 5 To remove multiple routers, holding the CTRL key, click the virtual routers that you want to delete. Right-click, and then click **Delete**.

Step 6 To save the changes, click **Save**.

The selected virtual routers are removed from the device configuration. All interfaces from the deleted virtual routers are automatically moved to the global virtual router, and all associated routing policies are permanently deleted.

Monitor virtual routers

To monitor and troubleshoot virtual routers, log into the device CLI and use these commands:

- **show vrf**: Displays the details of the virtual routers and their associated interfaces.
- **show route vrf** <*vrf_name*>: Displays the routing details of a virtual router.
- **show run router BGP all**: Displays the BGP routing details of all virtual router.
- **show run router BGP vrf** <*vrf_name*>: Displays the BGP routing details of a virtual router.
- **show crypto ipsec sa/show crypto ikev2 sa**: Displays the details of the tunnel and the associated virtual router.
- You can monitor the tunnels in the Site-to site monitoring dashboard (**Overview > Site to Site VPN**).

  In the **Tunnel Status** widget, hover over a topology, click View and click **Packet Tracer** to view and troubleshoot the threat defense VPN tunnels.

Configuration examples for virtual routers

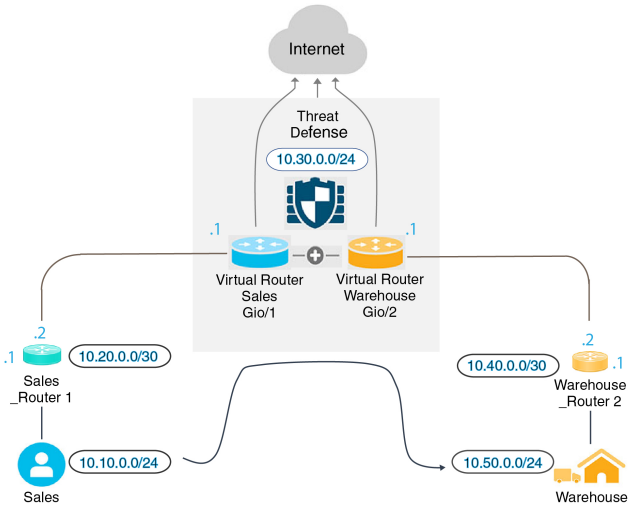
Route to a distant server through virtual routers

This task enables network separation while maintaining connectivity to servers across different virtual routing domains. You can create multiple virtual routers to maintain separate routing tables for groups of interfaces and interconnect them to reach hosts that are multiple hops away.

In virtual routing, you can create multiple virtual routers to maintain separate routing tables for groups of interfaces, thereby achieve network separation. In some scenarios, you may need to access a server that is reachable only through a separate virtual router. This example provides the procedure that interconnects virtual routers to reach to a host that is multiple hops away.

Consider an example, where a member of the sales department of a garment company wants to look up at the stock maintained by the warehousing department of its factory unit. In a virtual routing environment, you need to leak route between virtual routers where destination (warehousing department) is multiple hops away from sales department. This route leaking is done by adding multihop route leak, where, you configure a static route in Sales virtual router(source) to an interface in Warehouse virtual router (destination). As the destination network is multi-hop away, you also need to configure the Warehouse virtual router with the route to the destination network, namely 10.50.0.0/24.

Figure 326: Interconnecting two virtual routers - an example



Before you begin

This example assumes that you have already configured Sales_Router1 to route traffic from 10.20.0.1/30 interface to 10.50.0.5/24.

Follow these steps to route to a distant server through virtual routers:

Procedure

Step 1 Configure the inside interface (Gi0/1) of the device to be assigned to Sales virtual router.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, click the device name, and click **Interfaces**.

b) Edit the Gi0/1 interface.

- **Name**—For this example, VR-Sales.
- Select the **Enabled** checkbox.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Enter 10.30.0.1/24.

c) Click **Ok**.
d) Click **Save**.

Step 2 Configure the inside interface (Gi0/2) of the device to be assigned to Warehouse virtual router.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, click the device name, and click **Interfaces**.
b) Edit the Gi0/2 interface.
   - **Name**—For this example, VR-Warehouse.
   - Select the **Enabled** checkbox.
   - In **IPV4**, for **IP Type**, choose **Use Static IP**.
   - **IP Address**—Leave it blank. The system does not allow you to configure interfaces with same IP address (10.30.0.1/24), as you are yet to create user-defined virtual routers.
c) Click **Ok**.
d) Click **Save**.

Step 3 Create Sales and Warehouse virtual routers and assign their interfaces.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Manage Virtual Routers**.
c) Click **Add Virtual Router** and create Sales.
d) Click **Add Virtual Router** and create Warehouse.
e) Select Sales from virtual router drop-down, in **Virtual Router Properties**, add VR-Sales as **Selected Interface** and save.
f) Select Warehouse from virtual router drop-down, in **Virtual Router Properties**, add VR-Warehouse as **Selected Interface** and save.

Step 4 Revisit the VR-Warehouse interface configuration.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, click the device name, and click **Interfaces**.
b) Click **Edit** against VR-Warehouse interface. Specify the IP Address as 10.30.0.1/24. The system now allows you to configure with same IP address of VR-Sales, because the interfaces are seperately assigned to two different virtual routers.
c) Click **Ok**.
d) Click **Save**.

Step 5 Create network objects for the warehouse server—10.50.0.0/24, and for the warehouse gateway— 10.40.0.2/30.

a) In the Cloud-Delivered Firewall Management Center, choose **Objects** > **Object Management**.
b) Choose **Add Network** > **Add Object**.
   - **Name**—For this example, Warehouse-Server.
   - **Network**—Click Network and enter 10.50.0.0/24.
c) Click **Save**.
d) Choose **Add Network** > **Add Object**.
   - **Name**—For this example, Warehouse-Gateway.
   - **Network**—Click Host and enter 10.40.0.2.

e) Click **Save**.

Step 6 Define the route leak in Sales that points to the VR-Warehouse interface.

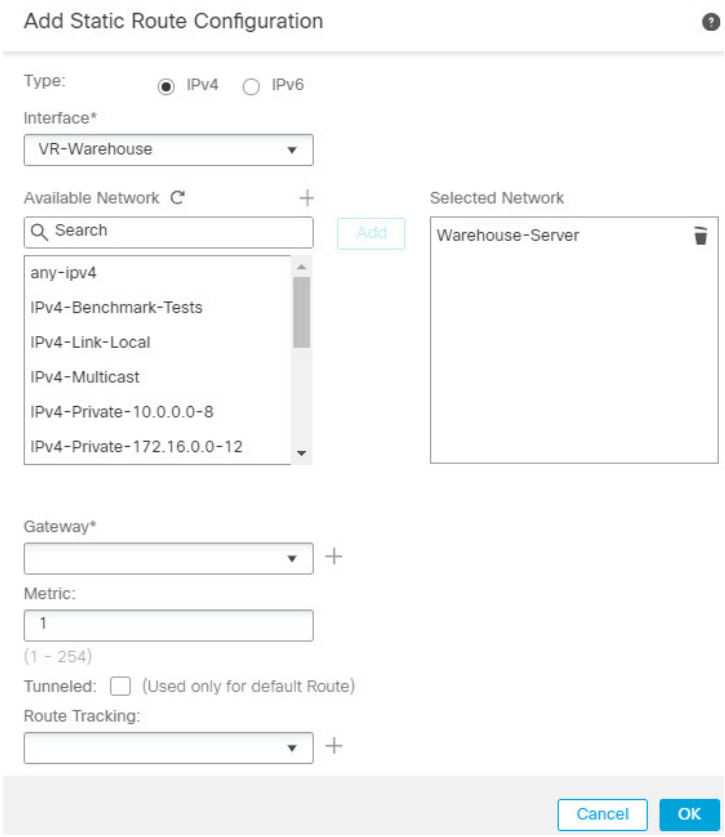
a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Choose **Routing**.

c) Choose Sales virtual router from the drop-down, and then click **Static Route**.

d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select VR-Warehouse.
- **Network**—Select the Warehouse-Server object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select the gateway.



e) Click **Ok**.

f) Click **Save**.

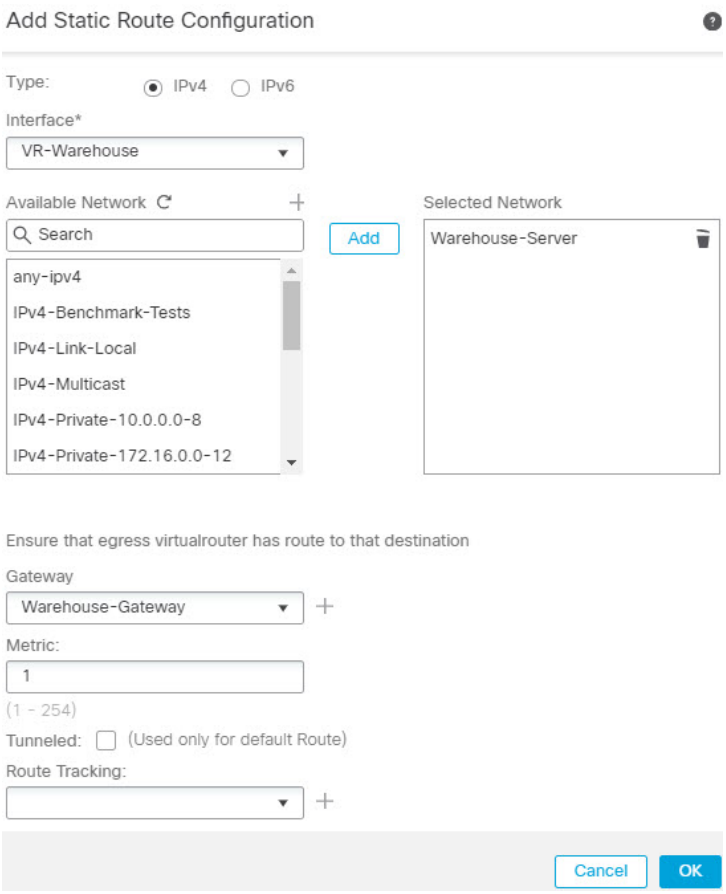
Step 7 In the Warehouse virtual router, define the route that points to the Warehouse Router 2 gateway.

a) In the Cloud-Delivered Firewall Management Center, choose Warehouse virtual router from the drop-down, and then click **Static Route**.

b) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select VR-Warehouse.
- **Network**—Select the Warehouse-Server object.

- **Gateway**—Select the Warehouse-Gateway object.

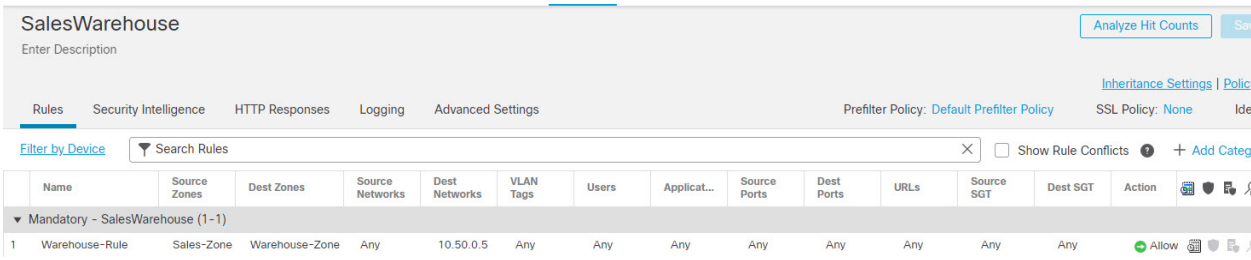


c) Click **Ok**.
d) Click **Save**.

Step 8 Configure access control rule that allows access to the warehouse server. For creating the access control rule, you need to create security zones. Use **Objects** > **Object Management** > **Interface**. Choose **Add** > **Security Zone** and create security zones for VR-Sales and VR-Warehouse; for Warehouse-Server network object, create a Warehouse-Server interface group (Choose **Add** > **Interface Group**).

Step 9 In the Cloud-Delivered Firewall Management Center, choose **Policies** > **Access Control heading** > **Access Control** and configure an access control rule to allow traffic from the source interfaces in the Sales virtual router to the destination interfaces in the Warehouse virtual router for the destination Warehouse-Server network object.

For example, if the interfaces in Sales are in the Sales-Zone security zone, and those in Warehouse are in the Warehouse-Zone security zone, the access control rule would look similar to this:



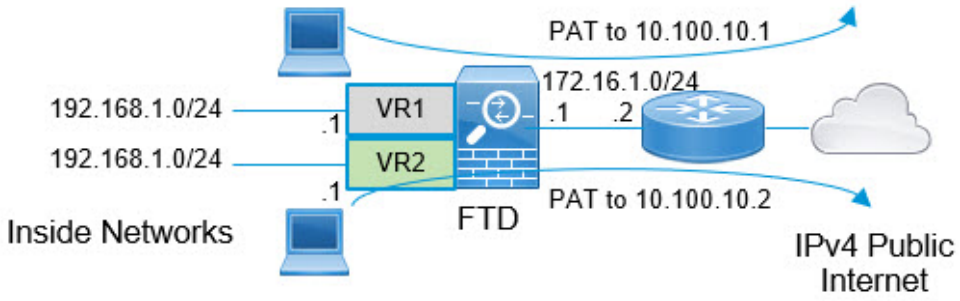
The route leak configuration enables traffic from the Sales virtual router to reach the warehouse server through the Warehouse virtual router. The access control rule allows the required traffic between the security zones.

How to Provide Internet Access with Overlapping Address Spaces

Ensure that endpoints using the same IP address space in different virtual routers can securely access the internet by implementing NAT/PAT policies

When using virtual routers, you can have the same network address for interfaces that reside in separate routers. However, because the IP addresses routed in these separate virtual routers are the same, apply NAT/PAT rules for each interface with separate NAT/PAT pools to ensure that return traffic goes to the correct destination. This example provides the procedure to configure the virtual routers and NAT/PAT rules to manage the overlapping address spaces.

For example, interfaces vr1-inside and vr2-inside on Firewall Threat Defense is defined to use the IP address 192.168.1.1/24, managing endpoints on their segment in the 192.168.1.0/24 network. To allow Internet access from two virtual routers that use the same address space, you need to apply NAT rules separately to the interfaces within each virtual router, ideally using separate NAT or PAT pools. You could use PAT to translate the source addresses in VR1 to 10.100.10.1, and for those in VR2, to 10.100.10.2. The following illustration shows this setup, where the Internet-facing outside interface is part of the global router. You must define the NAT/PAT rules with the source interface (vr1-inside and vr2-inside) explicitly selected—using “any” as the source interface makes it impossible for the system to identify the correct source because the same IP address could exist on two different interfaces.



Note Even if you have some interfaces within virtual routers that does not use overlapping address spaces, define the NAT rule with the source interface to make troubleshooting easier, and to ensure a cleaner separation between traffic from the virtual routers that is Internet-bound.

Procedure

Step 1 Configure the inside interface of the device for VR1:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, click the device name, and then click **Interfaces** .

b) Edit the interfaces that you want to assign to VR1:

- **Name**—For this example, vr1-inside.
- Select the **Enabled** checkbox.
- In **IPV4**, for **IP Type**, choose **Use Static IP** .
- **IP Address**—Enter 192.168.1.1/24.

c) Click **Ok** .

d) Click **Save** .

Step 2 Configure the inside interface of the device for VR2:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** , click the device name, and then click **Interfaces** .

b) Edit the interfaces that you want to assign to VR2:

- **Name** —For this example, vr2-inside.
- Select the **Enabled** checkbox.
- In **IPV4** , for **IP Type** , choose **Use Static IP** .
- **IP Address** —Leave it blank. The system does not allow you to configure interfaces with same IP address, as you are yet to create user-defined virtual routers.

c) Click **Ok** .

d) Click **Save** .

Step 3 Configure VR1 and the static default route leak to the outside interface:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** , and edit the Firewall Threat Defense device.

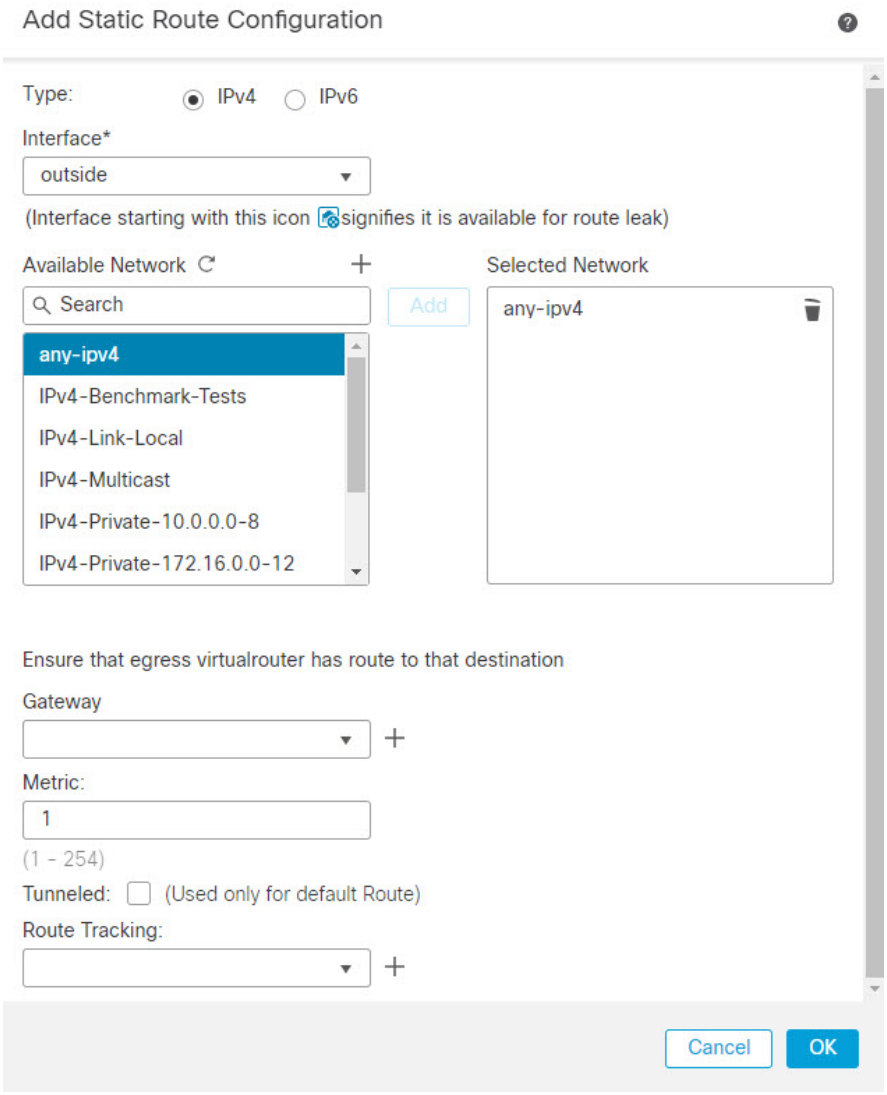
b) Choose **Routing** > **Manage Virtual Routers** . Click **Add Virtual Router** and create VR1.

c) For VR1, in **Virtual Router Properties** , assign vr1-inside and save.

d) Click **Static Route**.

e) Click **Add Route**. In **Add Static Route Configuration** , specify the following:

- **Interface** —Select the outside interface of the global router.
- **Network** —Select the any-ipv4 object. This network is the default route for any traffic that cannot be routed within VR1.
- **Gateway** —Leave it blank. When leaking a route into another virtual router, do not provide a Gateway.



f) Click **Ok**.

g) Click **Save**.

Step 4 Configure VR2 and the static default route leak to the outside interface:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** , and edit the Firewall Threat Defense device.

b) Choose **Routing** > **Manage Virtual Routers**. Click **Add Virtual Router** and create VR2.

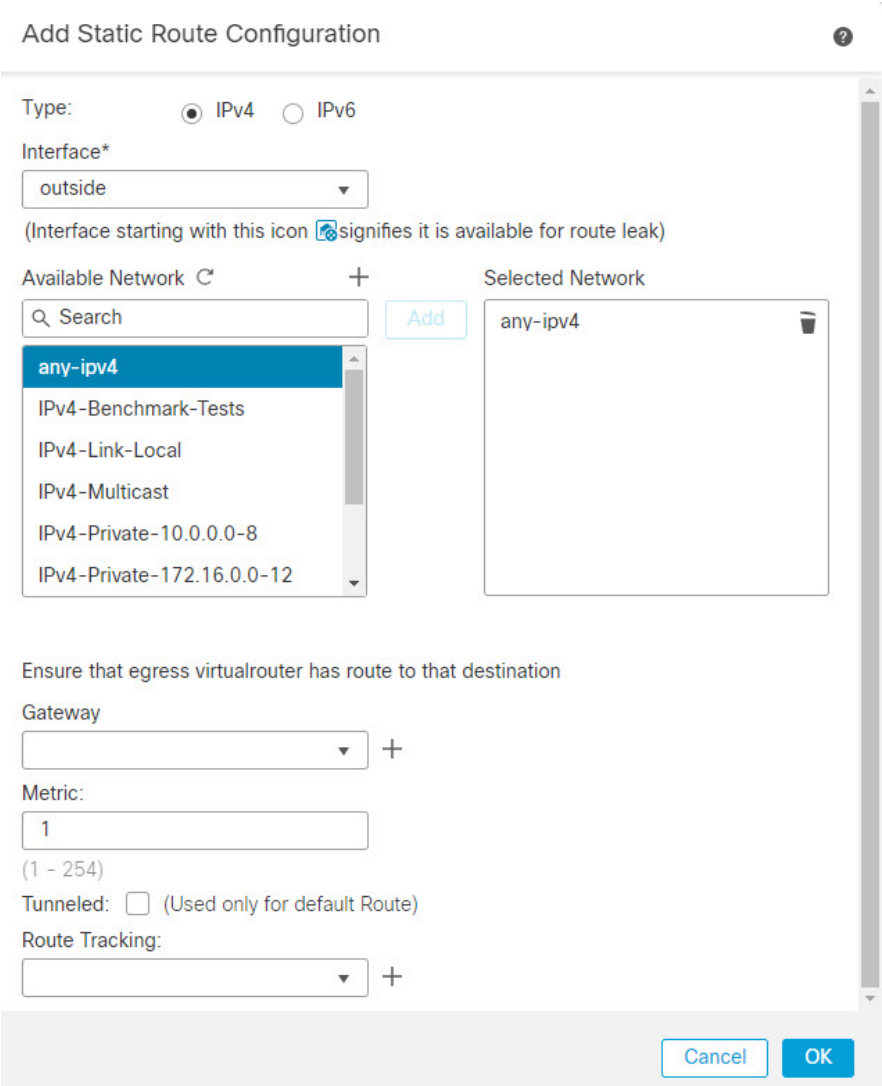
c) For VR2, in **Virtual Router Properties**, assign vr2-inside and save.

d) Click **Static Route**.

e) Click **Add Route**. In **Add Static Route Configuration** , specify the following:

- **Interface** —Select the outside interface of the global router.
- **Network** —Select the any-ipv4 object. This network is the default route for any traffic that cannot be routed within VR2.

- **Gateway** —Leave it blank. When leaking a route into another virtual router, do not select the Gateway.



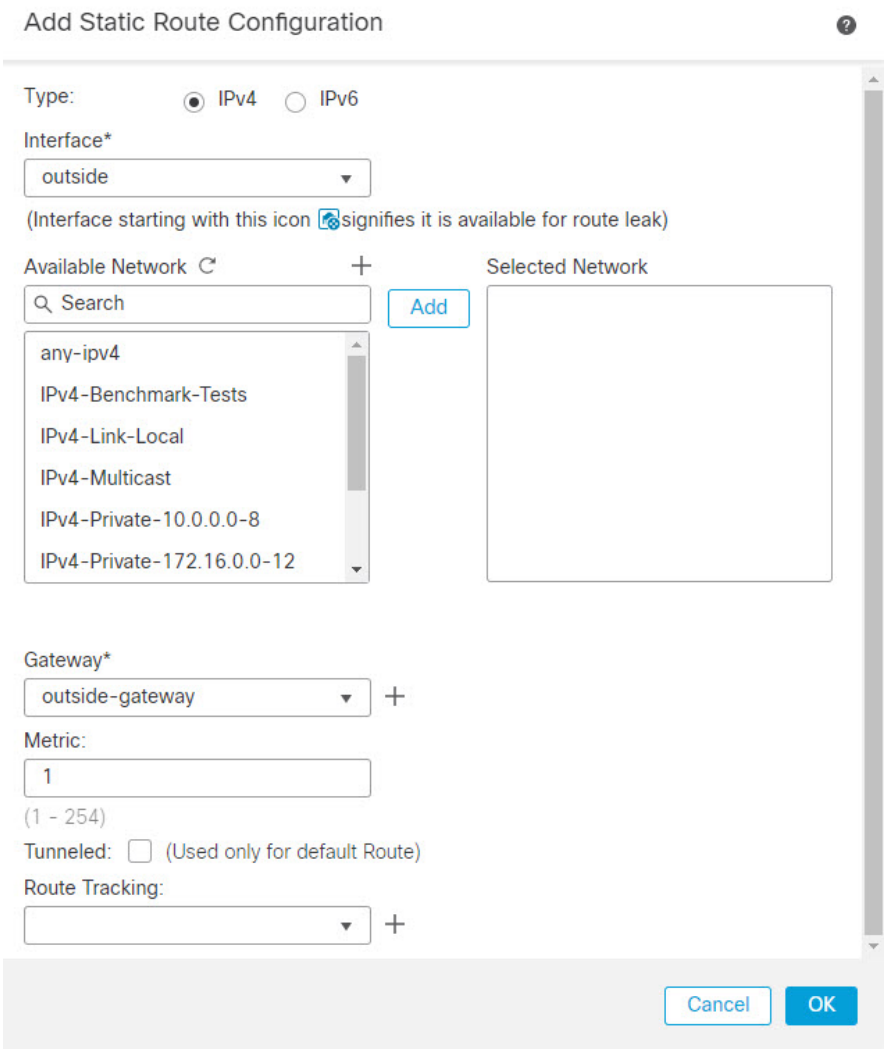
f) Click **Ok**.
g) Click **Save**.

Step 5 Configure IPv4 static default route, namely 172.16.1.2 on the outside interface of the global router:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** , and edit the Firewall Threat Defense device.
b) Choose **Routing** and edit global router properties.
c) Click **Static Route** .
d) Click **Add Route** . In **Add Static Route Configuration** , specify the following:

- **Interface** —Select the outside interface of the global router.
- **Network** —Select the any-ipv4 object. This will be the default route for any IPv4 traffic.

- **Gateway** —If already created, select the host name from the drop-down. If the object is not yet created, click **Add** and define the host object for the IP address of the gateway at the other end of the network link on the outside interface, in this example, 172.16.1.2. After you create the object, select it in the Gateway field.



e) Click **Ok** .
f) Click **Save** .

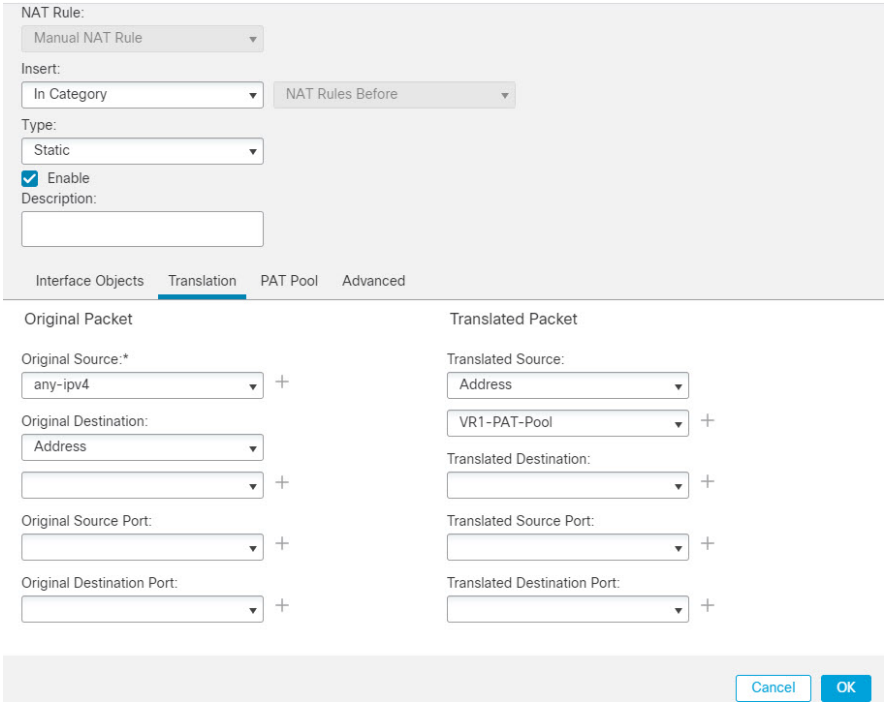
Step 6 Revisit the vr2-inside interface configuration:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** , click the device name, and then click **Interfaces** .
b) Click **Edit** against vr2-inside interface. Specify the IP Address as 192.168.1.1/24. The system now allows you to configure with same IP address of vr1-inside, because the interfaces are separately assigned to two different virtual routers.
c) Click **Ok** .
d) Click **Save** .

Step 7 Create the NAT rule to PAT inside to outside traffic of VR1 to 10.100.10.1.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **NAT**.
b) Click **New Policy** > **Threat Defense NAT**.
c) Enter InsideOutsideNATRule as the NAT policy name, and select the Firewall Threat Defense device. Click **Save** .
d) In InsideOutsideNATRule page, click **Add Rule** and define the following:

- **NAT Rule** —Select Manual NAT Rule.
- **Type** —Select Dynamic.
- **Insert** —Above, if any dynamic NAT rule exists.
- Click **Enabled** .
- In **Interface Objects** , select vr1-interface object and click **Add to Source** (If the object is not available, create one in **Objects** > **Object Management** > **Interface**), and select outside as **Add to Destination** .
- In **Translation** , for **Original Source** , select any-ipv4. For **Translated Source** , click **Add** and define host object VR1-PAT-Pool with 10.100.10.1. Select VR1-PAT-Pool as shown in the figure below:



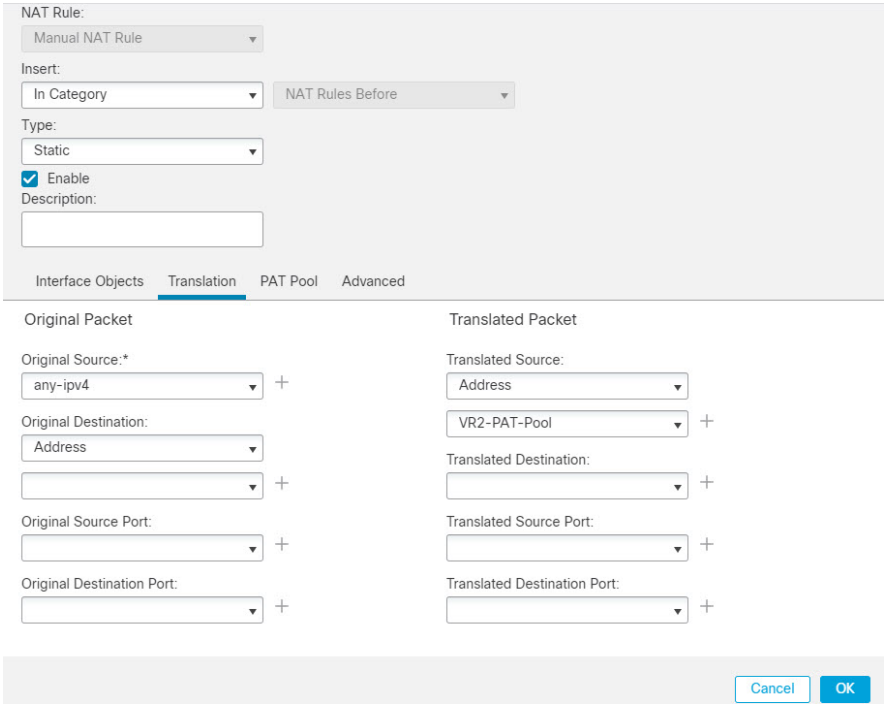
e) Click **Ok** .
f) Click **Save** .

Step 8 Add NAT rule to PAT inside to outside traffic of VR2 to 10.100.10.2.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **NAT** .
b) Edit InsideOutsideNATRule to define the VR2 NAT rule:

- **NAT Rule** —Select Manual NAT Rule.
- **Type** —Select Dynamic.

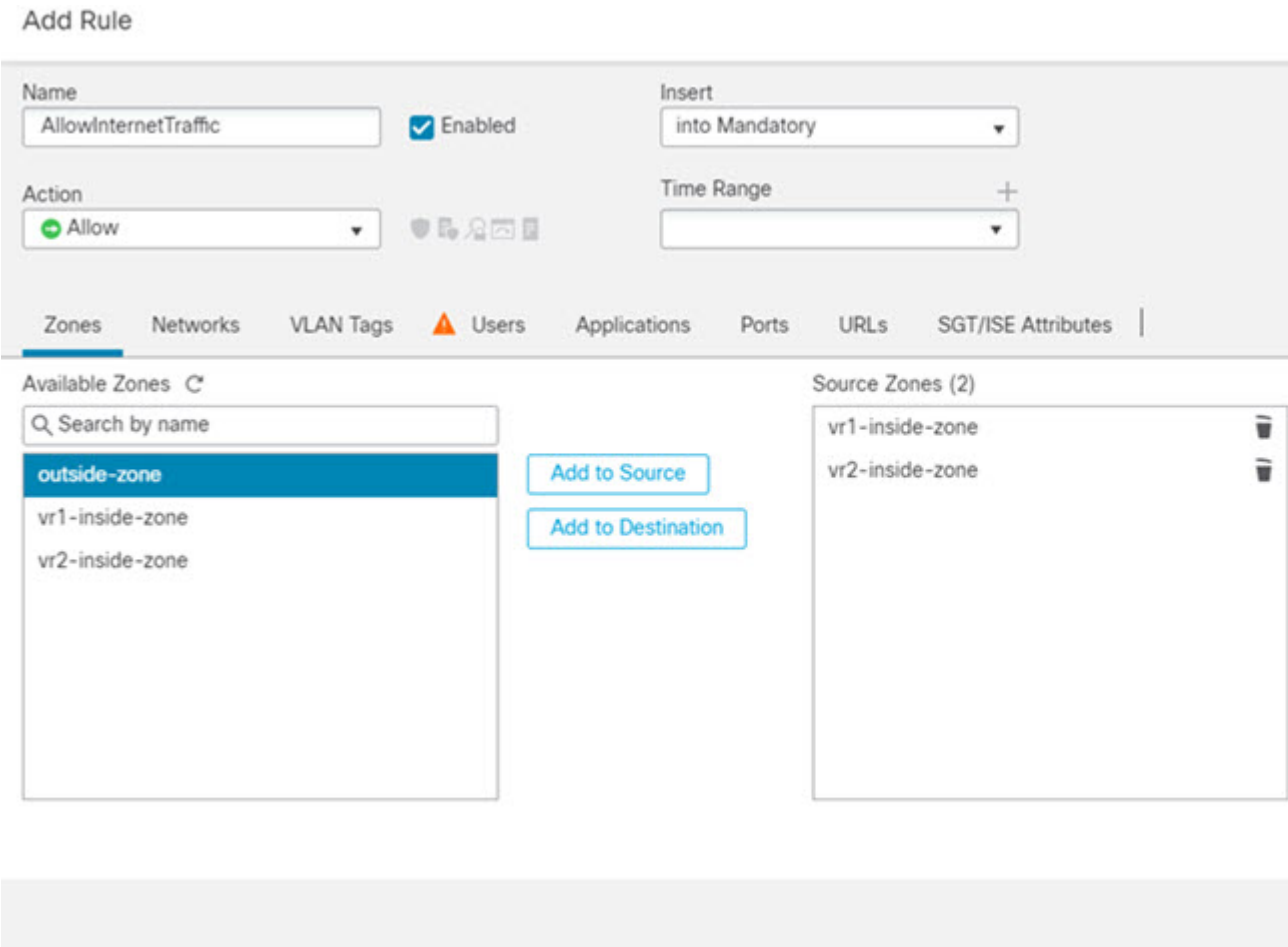
- **Insert** —Above, if any dynamic NAT rule exists.
- Click **Enabled** .
- In **Interface Objects** , select vr2-interface object and click **Add to Source** (If the object is not available, create one in **System** (⚙) > **Configuration** > **Management Interfaces**), and select outside as **Add to Destination** .
- In **Translation** , for **Original Source** , select any-ipv4. For **Translated Source** , click **Add** and define host object VR2-PAT-Pool with 10.100.10.2. Select VR2-PAT-Pool as shown in the figure below:



c) Click **Ok** .
d) Click **Save** .

Step 9 To configure the access control policy that allows traffic from the vr1-inside and vr2-inside interfaces to the outside interface, you need to create security zones. Use **Objects** > **Object Management** > **Interface** . Choose **Add** > **Security Zone** and create security zones for vr1-inside, vr2-inside, and outside interfaces. Choose **Policies** > **Access Control heading** > **Access Control** and configure an access control rule to allow traffic from vr1-inside-zone and vr2- inside-zone to outside-zone.

Assuming that you create zones named after the interfaces, a basic rule that allows all traffic to flow to the Internet will look like the following. You can apply other parameters to this access control policy:

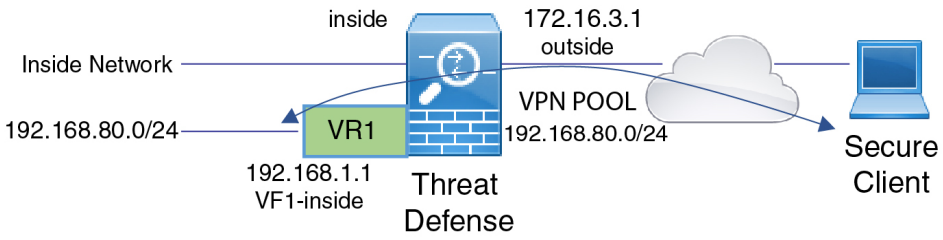


Allow RA VPN access to internal networks in virtual routing

Allow Secure Client users to connect to user-defined virtual router networks by configuring route leaks between the global virtual router and user-defined virtual routers.

On virtual routing-enabled devices, RA VPN is supported only on global virtual router interfaces. This example provides the procedure that allows your Secure Client user to connect to user-defined virtual router networks.

In the following example, the RA VPN (Secure Client) user connects to the outside interface of Firewall Threat Defense at 172.16.3.1, and is given an IP address within the pool of 192.168.80.0/24. The user can access the inside network of only the global virtual router. To allow traffic flow through the network of the user-defined virtual router VR1, namely 192.168.1.0/24, leak the route by configuring the static routes on global and VR1.



Before you begin

This example assumes that you have already configured the RA VPN, defined the virtual routers, and configured and assigned the interfaces to the appropriate virtual routers.

Follow these steps to allow RA VPN access to internal networks in virtual routing:

Procedure

Step 1 Configure route leak from Global virtual router to the user-defined VR1:

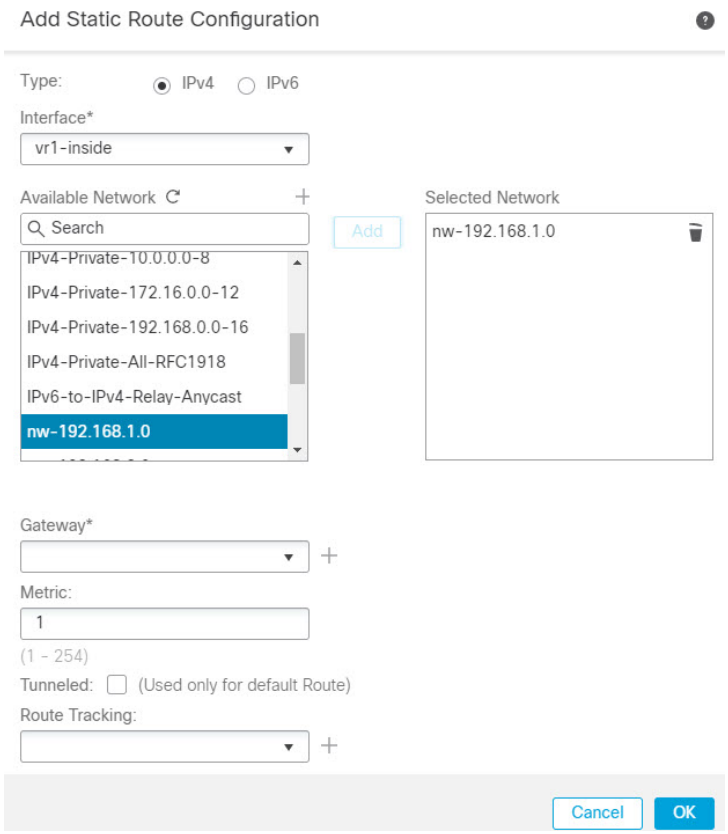
a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Click **Routing**. By default, the Global routing properties page appears.

c) Click **Static Route**.

d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the VR1 inside interface.
- **Network**—Select the VR1 virtual router network object. You can create one using the **Add Object** option.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, does not select the gateway.



The route leak allows Secure Client assigned IP addresses in the VPN pool to access the 192.168.1.0/24 network in the VR1 virtual router.

e) Click **Ok**.

Step 2 Configure the route leak from VR1 to the Global virtual router:

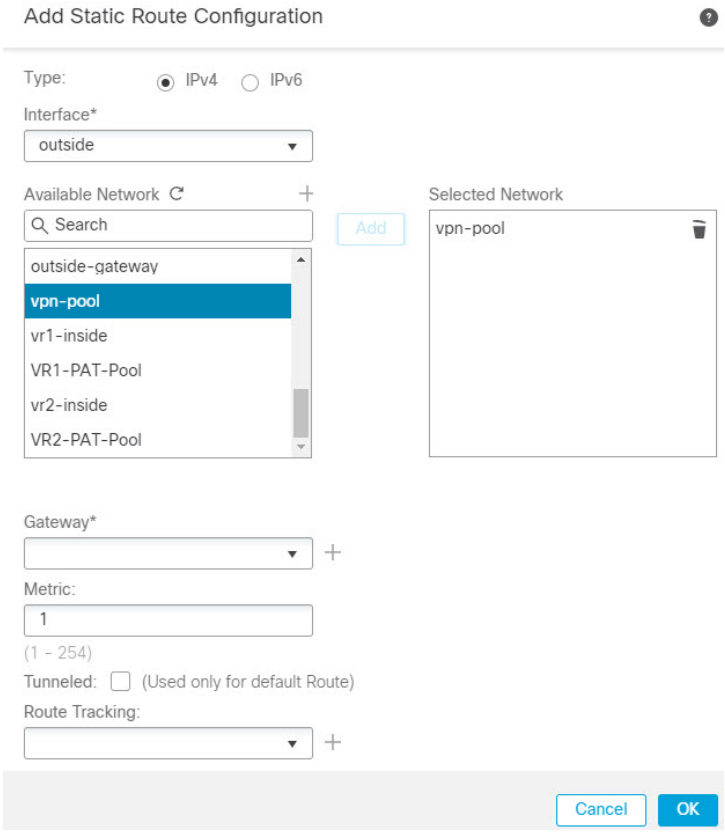
a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Click **Routing** and from the drop-down, select VR1.

c) Click **Static Route**.

d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the outside interface of the global router.
- **Network**—Select the global virtual router network object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, does not select the gateway.



The configured static route allows endpoints on the 192.168.1.0/24 network (VR1) to initiate connections to Secure Client assigned IP addresses in the VPN pool.

e) Click **Ok**.

Secure Client users can now access networks in both the global virtual router and the user-defined VR1 virtual router through the configured route leaks.

What to do next

If RA VPN address pool and the IP addresses in the user-defined virtual router overlap, you must also use static NAT rules on the IP addresses to enable proper routing. Alternatively, you can change your RA VPN address pool so that there is no overlap.

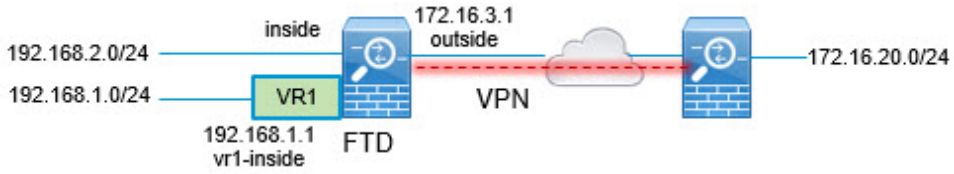
Secure traffic from networks in multiple virtual routers over a site-to-site VPN

This task allows you to secure connections from or to networks hosted within user-defined virtual routers over the site-to-site VPN by configuring route leaking between virtual routers and updating the VPN connection profile.

On virtual routing-enabled devices, Site-to-Site VPN is supported only on global virtual router interfaces. You cannot configure it on an interface that belongs to a user-defined virtual router. This procedure allows you to secure the connections from or to networks hosted within user-defined virtual routers over the site-to-site

VPN. You also need to update the site-to-site VPN connection to include the user-defined virtual routing networks.

Let us consider a scenario, where, a site-to-site VPN is configured between a branch office network to a company headquaters network; the Firewall Threat Defense in the branch office having virtual routers. In this case, the site-to-site VPN is defined on the outside interface of the branch office at 172.16.3.1. This VPN includes the inside network 192.168.2.0/24 without extra configuration, because the inside interface is also part of the global virtual router. But, to provide site-to-site VPN services to the 192.168.1.0/24 network, which is part of the VR1 virtual router, you must leak the route by configuring the static routes on global and VR1, and add the VR1 network to the site-to-site VPN configuration.



Before you begin

This example assumes that you have already configured the site-to-site VPN between the 192.168.2.0/24 local network and the 172.16.20.0/24 external network, defined the virtual routers, and configured and assigned the interfaces to the appropriate virtual routers.

Follow these steps to secure traffic from networks in multiple virtual routers over a site-to-site VPN:

Procedure

Step 1 Configure route leak from the Global virtual router to the user-defined VR1:

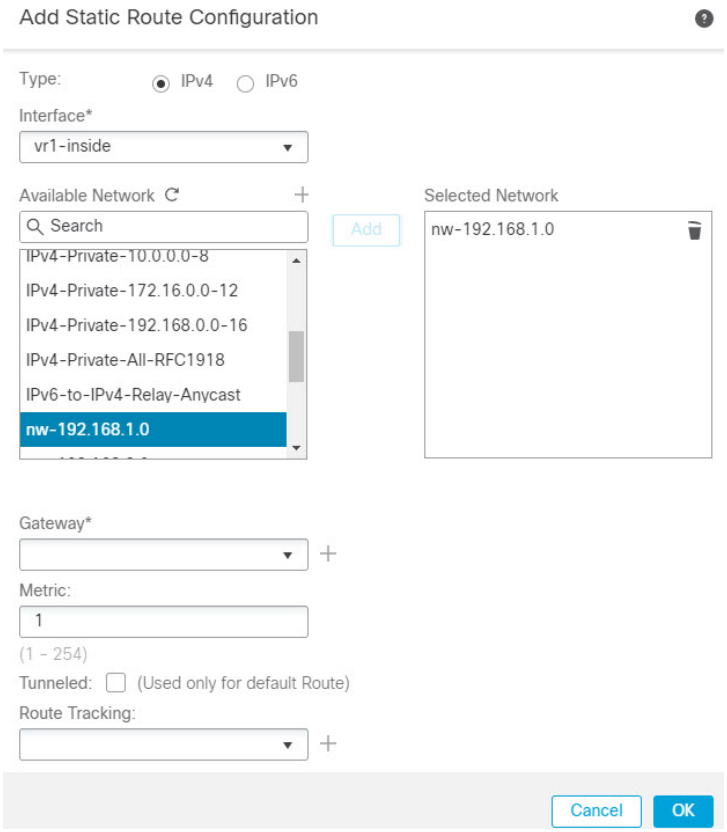
a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Click **Routing**. By default, the Global routing properties page appears.

c) Click **Static Route**.

d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the VR1 inside interface.
- **Network**—Select the VR1 virtual router network object. You can create one using the **Add Object** option.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select the gateway.



The route leak allows endpoints protected by the external (remote) end of the site-to-site VPN to access the 192.168.1.0/24 network in the VR1 virtual router.

e) Click **Ok**.

Step 2 Configure the route leak from VR1 to the Global virtual router:

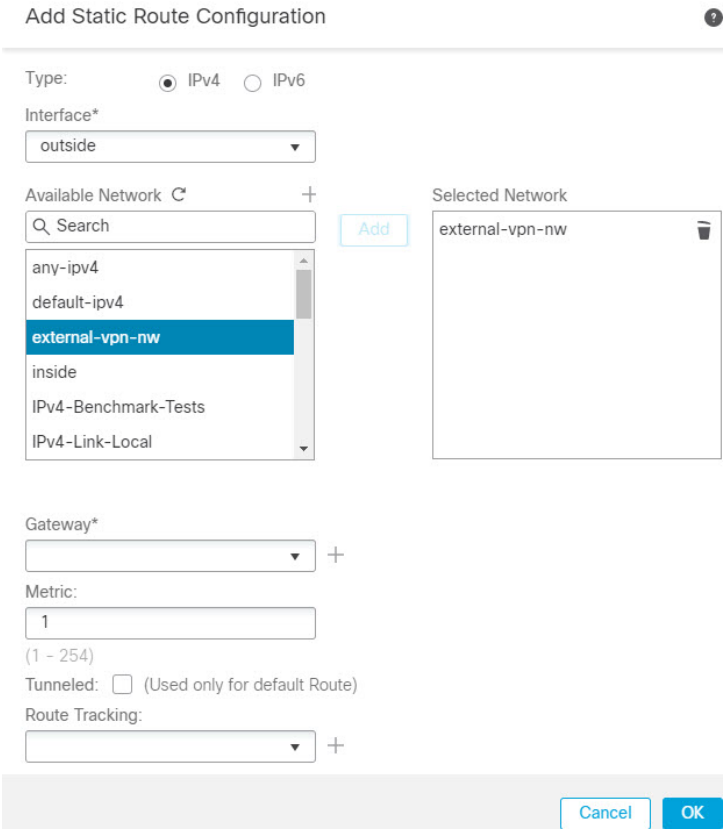
a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Click **Routing** and from the drop-down, select VR1.

c) Click **Static Route**.

d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the outside interface of the global router.
- **Network**—Select the global virtual router network object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select the gateway.



This static route allows endpoints on the 192.168.1.0/24 network (VR1) to initiate connections that will traverse the site-to-site VPN tunnel. For this example, the remote endpoint that is protecting the 172.16.20.0/24 network.

e) Click **Ok**.

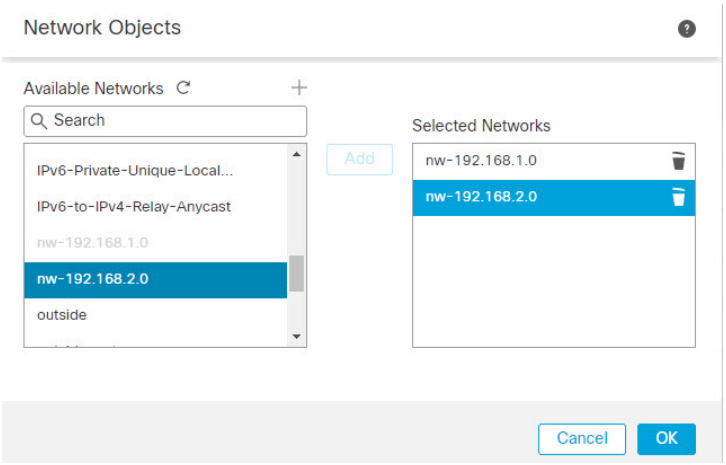
Step 3 Add the 192.168.1.0/24 network to the site-to-site VPN connection profile:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **VPN** > **Site to Site**, and edit the VPN Topology.

b) In **Endpoints**, edit Node A endpoint.

c) In **Edit Endpoint**, in the **Protected Networks** field, click **Add New Network Object**.

d) Add the VR1 network object with 192.168.1.0 network:



e) Click **Ok** and save the configuration.

The route leaking configuration allows traffic from networks in user-defined virtual routers to traverse the site-to-site VPN tunnel securely, enabling connectivity between the VR1 network and remote networks protected by the VPN.

Secure traffic from networks with multiple virtual routers over a site-to-site VPN with dynamic VTI

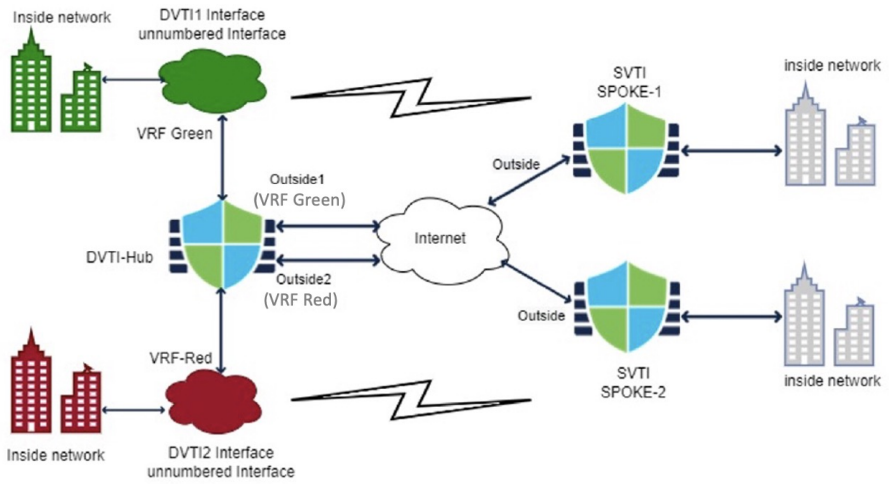
This task enables ISPs to create virtual routers and associate dynamic VTIs with these virtual routers to extend the capabilities of dynamic VTIs in their network. You can associate dynamic VTIs either with global or user-defined virtual routers.

ISPs have different segmented networks for different customers. You can create virtual routers, associate dynamic VTIs with these virtual routers, and extend the capabilities of dynamic VTIs in your network. You can associate dynamic VTIs either with global or user-defined virtual routers. A single threat defense device can act as a dynamic VTI hub with a global or one or more user-defined virtual routers. Each user-defined virtual router can be one customer network.

Let us consider an example where route-based site-to-site VPNs are configured between two company headquarter networks and their two branch office networks. The ISP's threat defense, a dynamic VTI hub, manages the two company headquarter networks with two user-defined virtual routers: VRF green and VRF red. The dynamic VTI hub establishes a site-to-site VPN between:

- Customer 1 (VRF green) and Branch 1 (SVTI spoke 1)
- Customer 2 (VRF red) and Branch 2 (SVT2 spoke 2)

Figure 327: Site-to-Site VPNs with multiple virtual routers and dynamic VTI



Procedure

Step 1 Configure a dynamic VTI interface on the hub.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and edit the threat defense device.
b) Choose **Add Interfaces > Virtual Tunnel Interface**.
c) Select the **Tunnel Type** as **Dynamic**.
d) Specify the interface name as DVTI1 and configure all the parameters for the dynamic VTI.
e) Click **Save**
f) Repeat steps 1a - e to configure the second dynamic VTI on the hub, DVTI2.

Step 2 Configure static VTI on spoke 1.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and edit the threat defense device.
b) Select **Add Interfaces > Virtual Tunnel Interface**.
c) Select the **Tunnel Type** as **Static**.
d) Specify the interface name as SVTI spoke-1 and configure all the parameters for the static VTI.
e) Click **Save**
f) Repeat steps 2a - e to configure the static VTI on spoke 2: SVTI spoke-2.

Step 3 Configure a route-based site-to-site VPN between hub and SVTI spoke 1.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **VPN** > **Site to Site** and click **+ Site To Site VPN**.
b) Enter a name for the VPN topology in the **Topology Name** field.
c) Choose **Route Based (VTI)** and select **Hub and Spoke** as the network topology.
d) Click the **Endpoints** tab.
e) Configure the hub and the spoke (DVTI1 and SVTI spoke-1) and their routing policies.
f) Configure the **IKE**, **IPsec**, and **Advanced** options for the VPN, if required.
g) Click **Save**.

h) Repeat steps 3a - g to configure the second route-based site-to-site VPN topology between hub (DVTI2) and SVTI spoke 2.

Step 4 Configure the two virtual routers.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and edit the threat defense device.
b) Click **Routing**.
c) Click **Manage Virtual Routers**.
d) Click **Add Virtual Router**.

Specify the name as VRF green and provide a description for the virtual router.

e) Repeat steps 4a - d to configure VRF red.

Step 5 Assign all the interfaces to the virtual routers.

a) From the drop-down list, select the virtual router.
b) In the **Virtual Router Properties** page, select the interfaces listed under the **Available Interfaces** box.

Assign the dynamic VTI interface along with the other interfaces.

c) Click **Add**.

Step 6 Repeat steps 5a - c for VRF red.

Step 7 Configure the routing policy for the virtual router.

a) From the drop-down list, select the virtual router.
b) Click **Static Route** or one of the dynamic routing protocols.
c) Configure the routing parameters.
d) Click **Save**.

You have successfully configured networks with multiple virtual routers over a site-to-site VPN with dynamic VTI, enabling ISPs to manage different segmented networks for different customers through virtual routers.

What to do next

Select the hub and spoke devices and click **Deploy**. After the deployment, you can monitor the VPN tunnels in the Site-to site monitoring dashboard (**Overview** > **Dashboards** > **Site to Site VPN**).

You can also use the commands listed in Monitor virtual routers, on page 1390 to view and troubleshoot the virtual router.

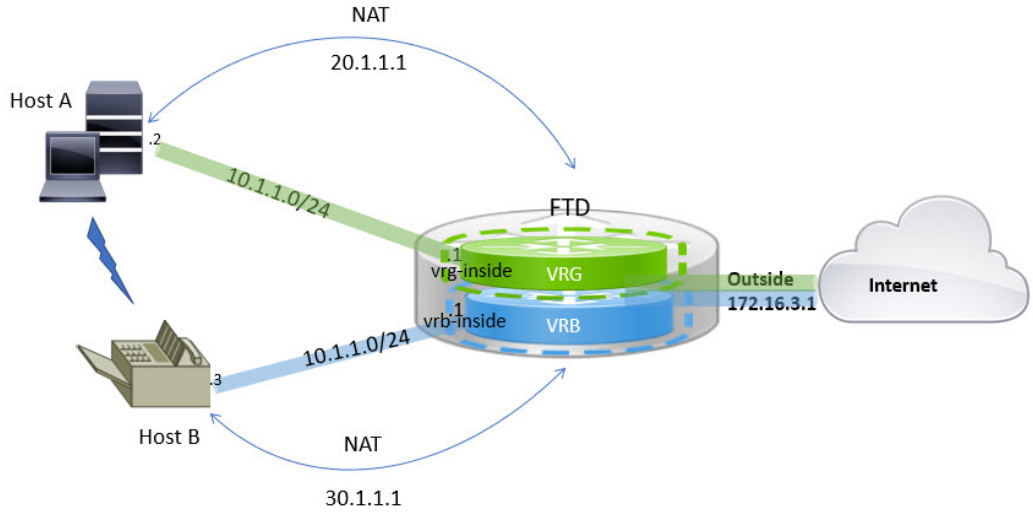
Route traffic between two overlapping network hosts in virtual routing

Configure NAT rules to manage overlapping network hosts and enable communication between hosts on different virtual routers that have the same network address.

You can configure hosts on the virtual routers that have same network address. If the hosts want to communicate, you can configure twice NAT. This example provides the procedure to configure the NAT rules to manage the overlapping network host.

In the following example, two hosts Host A and Host B belong to different virtual routers: VRG (interface VRG-inside), VRB (interface VRB-inside) respectively with the same subnet 10.1.1.0/24. For both the hosts to communicate, create a NAT policy where, VRG-Host interface object would use a mapped NAT address

- 20.1.1.1, and VRB-Host interface object would use a mapped NAT address - 30.1.1.1. Thus, Host A uses 30.1.1.1 to communicate to Host B; Host B uses 20.1.1.1 to reach Host A.



Before you begin

This example assumes that you have already configured:

- VRG-inside and VRB-inside interfaces are associated with virtual routers: VRG and VRB respectively and VRG-inside and VRB-inside interfaces configured with same subnet address (say, 10.1.1.0/24).
- Interfaces zones VRG-Inf, VRB-Inf created with VRG-inside and VRB-inside interfaces respectively.
- Host A in VRG with VRG-inside as default gateway; Host B in VRB with VRB-inside as default gateway.

Follow these steps to route traffic between two overlapping network hosts in virtual routing:

Procedure

Step 1 Create the NAT rule to handle traffic from Host A to Host B. In the Cloud-Delivered Firewall Management Center, choose **Devices** > **NAT**.

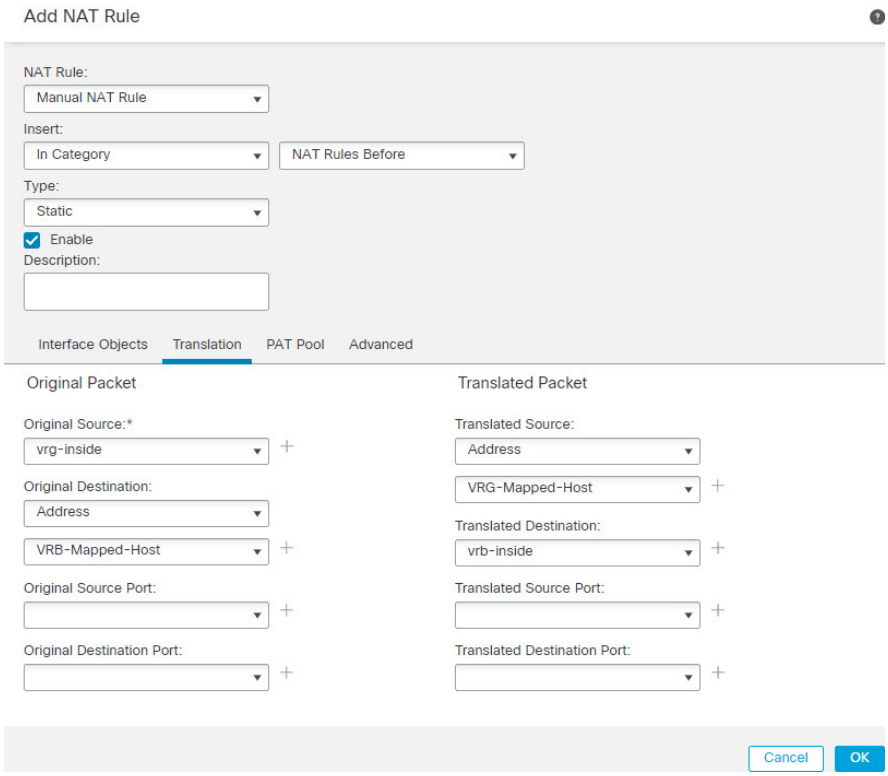
Step 2 Click **New Policy** > **Threat Defense NAT**.

Step 3 Enter a NAT policy name, and select the Firewall Threat Defense device. Click **Save**.

Step 4 In the NAT page, click **Add Rule** and define the following:

- **NAT Rule**—Select Manual NAT Rule.
- **Type**—Select Static.
- **Insert**—Select Above, if any NAT rule exists.
- Click **Enabled**.
- In **Interface Objects**, select VRG-Inf object and click **Add to Source** (If the object isn't available, create one in **Objects** > **Object Management** > **Interface**), and select VRB-Inf object and click **Add to Destination**.
- In **Translation**, select the following:

- **Original Source**, select VRG-inside.
- **Original Destination**, click **Add** and define object VRB-Mapped-Host with 30.1.1.1. Select VRB-Mapped-Host.
- **Translated Source**, click **Add** and define object, VRG-Mapped-Host with 20.1.1.1. Select VRG-Mapped-Host.
- **Translated Destination**, select VRB-inside as shown in the following figure:



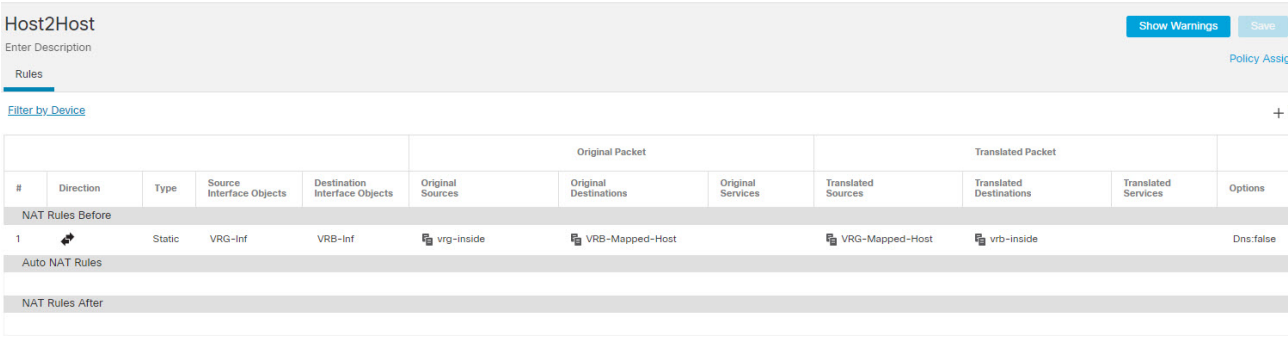
When you run the **show NAT detail** command on the Firewall Threat Defense device, you will see an output similar to this:

```
firepower(config-service-object-group)# show nat detail
Manual NAT Policies (Section 1)
1 (2001) to (3001) source static vrg-inside VRG-MAPPED-HOST destination static VRB-MAPPED-HOST
 vrb-inside
translate_hits = 0, untranslate_hits = 0
Source - Origin: 10.1.1.1/24, Translated: 20.1.1.1/24
Destination - Origin: 30.1.1.1/24, Translated: 10.1.1.1/24
```

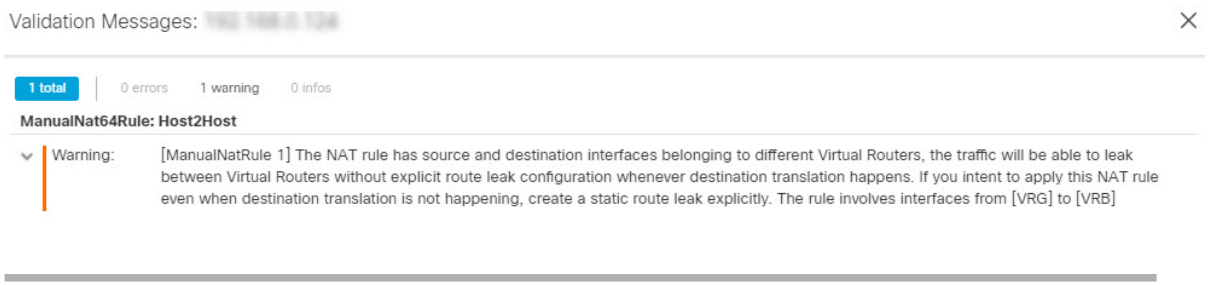
Step 5 Click **Ok**.

Step 6 Click **Save**.

The NAT rule looks like this:



When you deploy the configuration, a warning message appears:



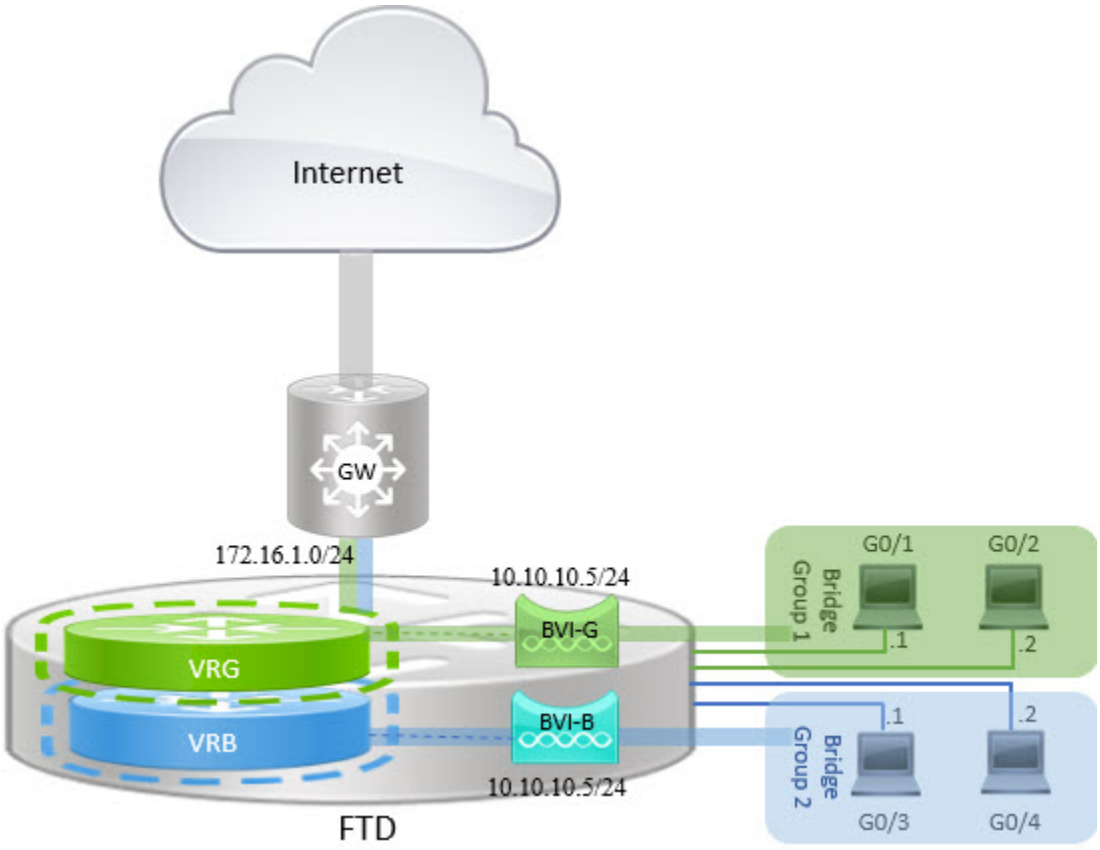
The NAT rule is configured to enable communication between Host A and Host B on overlapping network addresses. Host A uses 30.1.1.1 to communicate with Host B, and Host B uses 20.1.1.1 to reach Host A.

Manage overlapping segments in routed firewall mode with BVI interfaces

This task allows you to deploy single Firewall Threat Defense between multiple overlapping networks transparently and/or deploy the firewall between the hosts of same network by configuring BVI per virtual router.

BVI is a virtual interface within a router that acts like a normal routed interface. It does not support bridging, but represents the comparable bridge group to routed interfaces within the router. All the packets coming in or going out of these bridged interfaces, pass through the BVI interface. The interface number of the BVI is the number of the bridge group that the virtual interface represents.

In the following example, BVI-G is configured in VRG and Bridge Group 1 is the routed interface for interfaces G0/1 and G0/2. Similarly, BVI-B is configured in VRB and Bridge Group 2 is the routed interface for interfaces G0/3 and G0/4. Consider that both BVIs have the same IP subnet address, say 10.10.10.5/24. Because of virtual routers, the network is isolated on the shared resources.



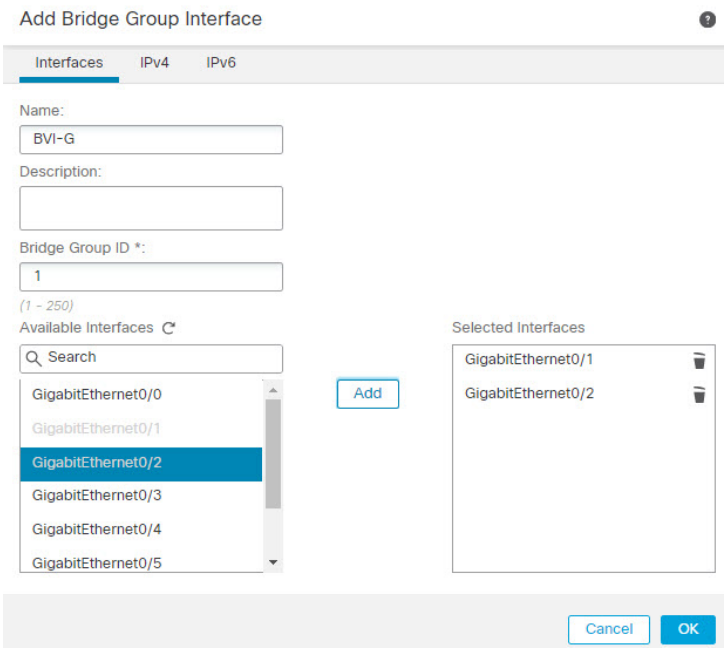
Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**. Edit the required device.

Step 2 In **Interfaces**, choose **Add Interfaces** > **Bridge Group Interface**.

a) Enter the following details for BVI-G:

- **Name**—For this example, BVI-G.
- **Bridge Group ID**—For this example, 1.
- **Available Interface**—Select the interfaces.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Enter 10.10.10.5/24.

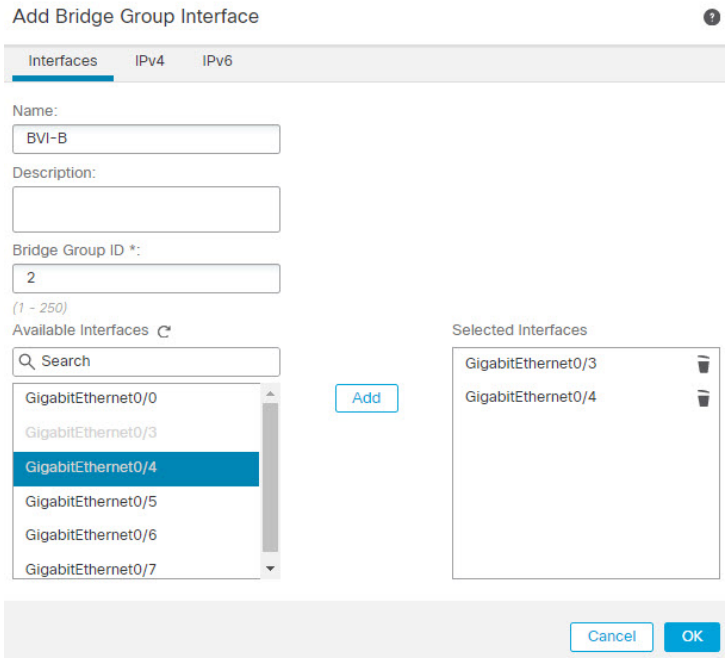


b) Click **Ok**.

c) Click **Save**.

a) Enter the following details for BVI-B:

- **Name**—For this example, BVI-B.
- **Bridge Group ID**—For this example, 2.
- **Available Interface**—Select the sub interfaces.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Leave this field empty as the system does not allow two interfaces to have overlapping IP address. You can revisit the Bridge Group and provide the same IP address after aligning it under a virtual router.



b) Click **Ok**.

c) Click **Save**.

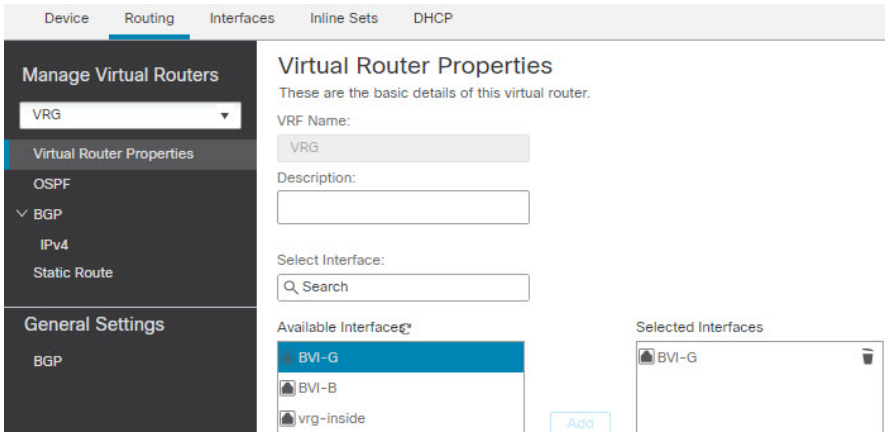
Step 3 Create virtual router, say VRG, and select BVI-G as its network:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**.

b) Edit the device, and choose **Routing** > **Manage Virtual Routers**.

c) Click **Add Virtual Router**. Enter a name for the virtual router and click **Ok**.

d) In **Virtual Routing Properties**, select **BVI-G** and click **Add**.



e) Click **Save**.

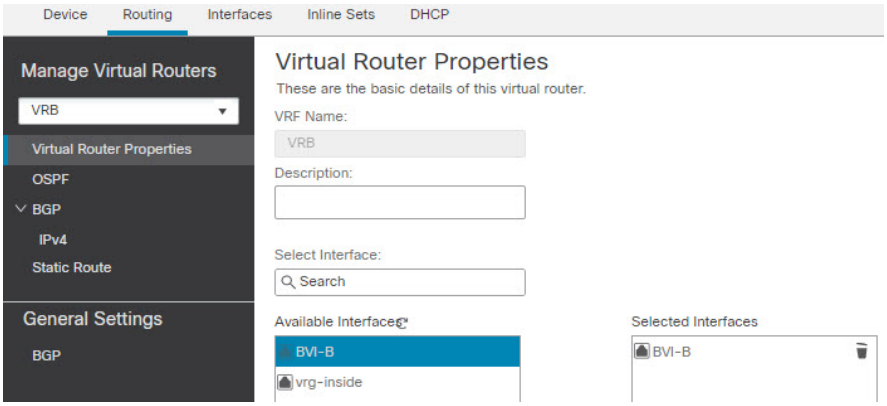
Step 4 Create virtual router, say VRB, and select BVI-B as its network:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**.

b) Edit the device, and choose **Routing** > **Manage Virtual Routers**.

c) Click **Add Virtual Router**. Enter a name for the virtual router and click **Ok**.

d) In **Virtual Routing Properties**, select **BVI-B** and click **Add**.



e) Click **Save**.

Step 5 Revisit the BVI-B configuration:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, edit the device and click **Interfaces**.

b) Click **Edit** against BVI-B interface. Specify the IP Address as 10.10.10.5/24. The system now allows you to configure with same IP address of BVI-G, because the interfaces are seperately assigned to two different virtual routers.

c) Click **Ok**.

d) Click **Save**.

If you want to enable inter-BVI communication, use an external router as default gateway. In overlapping BVI scenarios, as in this example, use twice NAT external router as gateway to establish inter-BVI traffic. When configuring NAT for the members of a bridge group, you specify the member interface. You cannot configure NAT for the bridge group interface (BVI) itself. When doing NAT between bridge group member interfaces, you must specify the real and mapped addresses. You cannot specify "any" as the interface.

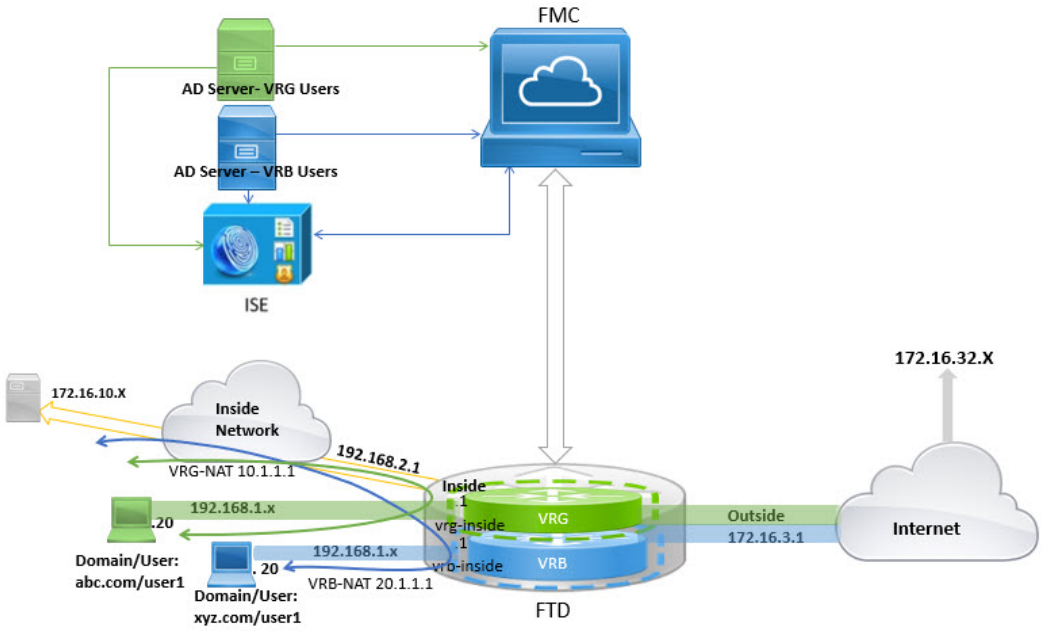
You have successfully configured BVI interfaces per virtual router to manage overlapping segments in routed firewall mode. The virtual routers isolate the network on shared resources, allowing the use of same IP subnet addresses for both BVIs.

Configure user authentication with overlapping networks

This task enables you to configure user authentication in virtual routing environments where multiple virtual routers share overlapping IP addresses and user networks. The configuration allows users from different domains to access shared resources while maintaining proper access control.

In virtual routing, you can configure multiple virtual routers with overlapping IP and overlapping users. In the example, VRG, and VRB are the virtual routers with overlapping IP - 192.168.1.1/24. The users on two different domains are also on overlapping network IP 192.168.1.20. For VRG and VRB users to access the shared server 172.16.10.X, leak routes to the global virtual router. Use source NAT to handle the overlapping IP. For controlling the access from VRG and VRB users, you must set user authentication in Cloud-Delivered Firewall Management Center. Cloud-Delivered Firewall Management Center uses realms, Active Directories, Identity source, and Identity rules and policies for authenticating user identity. Because Firewall Threat Defense

does not have direct role in authenticating users, user access is managed only through the access control policy. For controlling traffic from the overlapping users, use Identity policy and rules to create access control policy.



Before you begin

This example assumes that you have:

- Two AD servers for the VRG and VRB users.
- ISE with the two AD servers added.

Follow these steps to configure user authentication with overlapping networks:

Procedure

Step 1 Configure the inside interface of the device for VRG:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, click the device name, and then click **Interfaces**.

b) Edit the interfaces that you want to assign to VRG:

- **Name**—For this example, VRG-inside.
- Select the **Enabled** checkbox.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Enter 192.168.1.1/24.

c) Click **Ok**.

d) Click **Save**.

Step 2 Configure the inside interface of the device for VRB:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** > **Interfaces**.

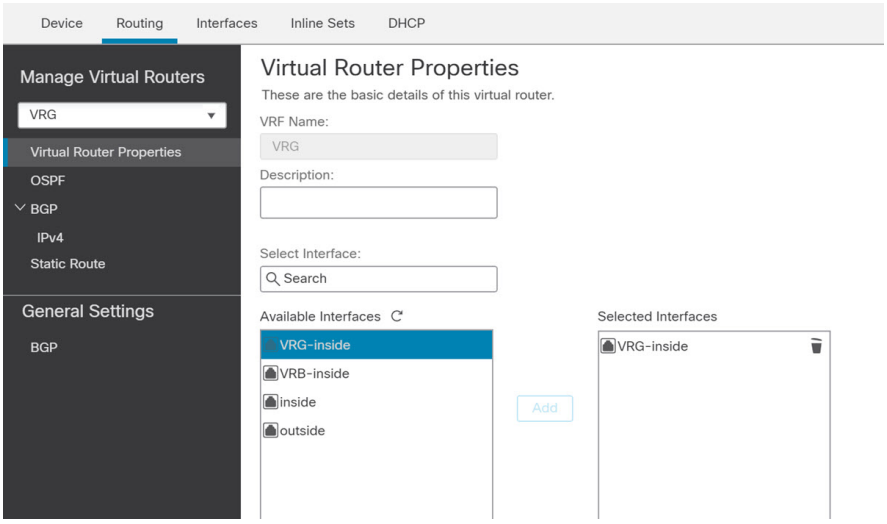
b) Edit the interfaces that you want to assign to VRB:

- **Name**—For this example, VRB-inside.
- Select the **Enabled** checkbox.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Leave it blank. The system doesn't allow you to configure interfaces with same IP address, as you're yet to create user-defined virtual routers.

c) Click **Ok**.
d) Click **Save**.

Step 3 Configure VRG and the static default route leak to the inside interface of the Global router for the VRG users to access the common server 172.16.10.1:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Manage Virtual Routers**. Click **Add Virtual Router** and create VRG.
c) For VRG, in **Virtual Router Properties**, assign VRG-inside and save.



d) Click **Static Route**.
e) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

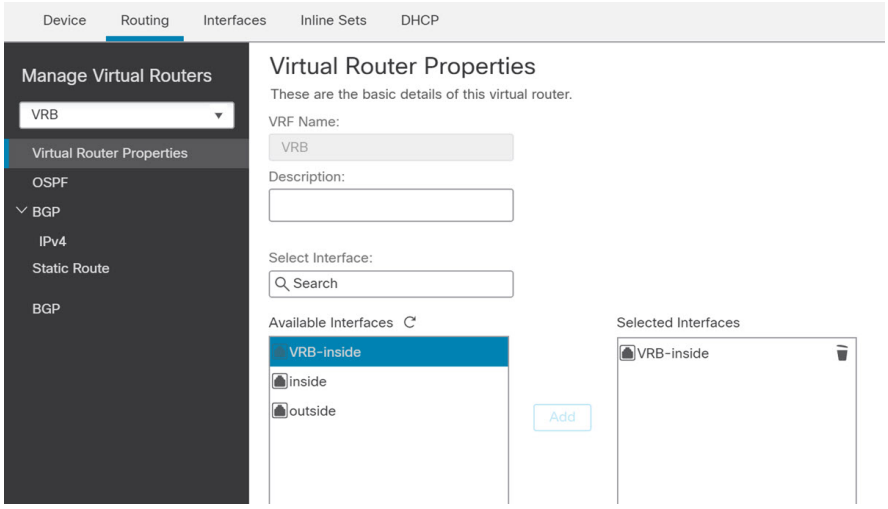
- **Interface**—Select the inside interface of the global router.
- **Network**—Select the any-ipv4 object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select a gateway.

f) Click **Ok**.
g) Click **Save**.

Step 4 Configure VRB and the static default route leak to the inside interface of the Global router for the VRB users to access the shared server 172.16.10.x:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Manage Virtual Routers**. Click **Add Virtual Router** and create VRB.

c) For VRB, in **Virtual Router Properties**, assign VRB-inside and save.



d) Click **Static Route**.

e) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the inside interface of the global router.
- **Network**—Select the any-ipv4 object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select a gateway.

f) Click **Ok**.

g) Click **Save**.

Step 5 Revisit the VRB-inside interface configuration:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, click the device name, and then click **Interfaces**.

b) Click **Edit** against VRB-inside interface. Specify the IP Address as 192.168.1.1/24. The system now allows you to configure with the same IP address as that of VRG-inside, because the interfaces are seperately assigned to two different virtual routers.

c) Click **Ok**.

d) Click **Save**.

Step 6 Add NAT rules for the source objects VRG and VRB. Click **Devices** > **NAT**.

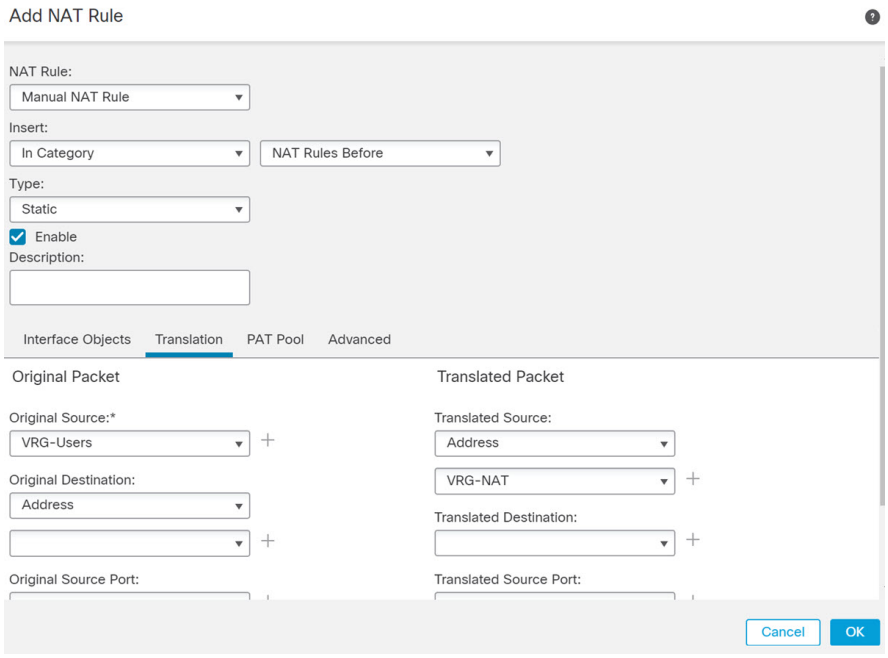
Step 7 Click **New Policy** > **Threat Defense NAT**.

Step 8 Enter a NAT policy name, and select the Firewall Threat Defense device. Click **Save**.

Step 9 In the NAT page, click **Add Rule** and define the following source NAT for VRG:

- **NAT Rule**—Select Manual NAT Rule.
- **Type**—Select Static.
- **Insert**—Select Above, if any NAT rule exists.
- Click **Enabled**.

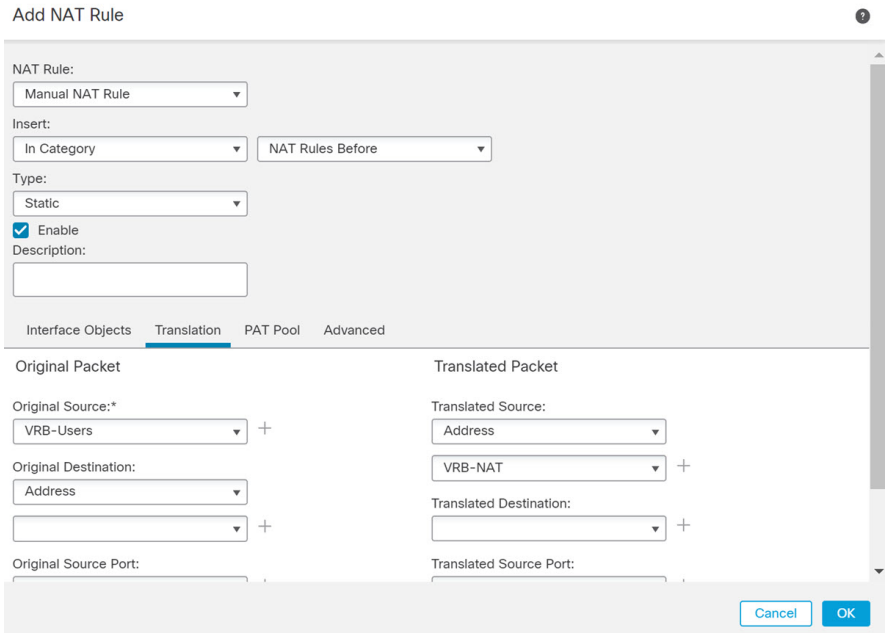
- In **Interface Objects**, select VRG-Inside object and click **Add to Source** (If the object is not available, create one in **Objects** > **Object Management** > **Interface**), and select Global-Inside object and click **Add to Destination**.
- In **Translation**, select the following:
  - **Original Source**, select VRG-Users.
  - **Translated Source**, click **Add** and define object, VRG-NAT with 10.1.1.1. Select VRG-NAT as shown in the figure:



Step 10 Click **Ok**.

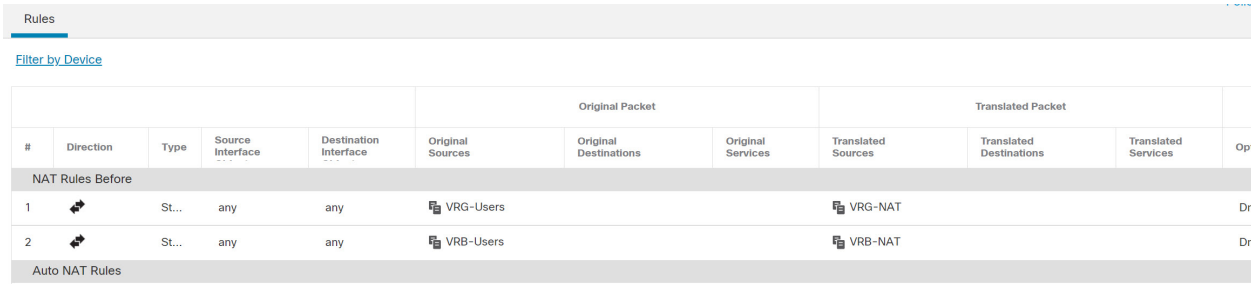
Step 11 In the NAT page, click **Add Rule** and define the following source NAT for VRB:

- **NAT Rule**—Select Manual NAT Rule.
- **Type**—Select Static.
- **Insert**—Select Above, if any NAT rule exists.
- Click **Enabled**.
- In **Interface Objects**, select VRB-Inside object and click **Add to Source** (If the object is not available, create one in **Objects** > **Object Management** > **Interface**), and select Global-Inside object and click **Add to Destination**.
- In **Translation**, select the following:
  - **Original Source**, select VRB-Users.
  - **Translated Source**, click **Add** and define object, VRB-NAT with 20.1.1.1. Select VRB-NAT as shown in the figure:



Step 12 Click **Save**.

The NAT rule looks like this:



Step 13 Add the two unique AD servers in Cloud-Delivered Firewall Management Center one for each VRG and VRB users—choose **System** > **Integration** > **Realms**.

Step 14 Click **New Realm** and complete the fields. For detailed information on the fields, see Realm fields, on page 2281.

Step 15 For controlling the access from VRG and VRB users, define 2 Active Directories, see Realm directory and synchronize fields, on page 2285see Create an LDAP realm or an Active Directory realm and realm directory, on page 2325

Step 16 Add ISE in Cloud-Delivered Firewall Management Center—choose **System** > **Integration** > **Identity Sources**.

Step 17 Click **Identity Services Engine** and complete the fields. For detailed information on the fields, see Configure ISE/ISE-PIC for user control using a realm, on page 2380.

Step 18 Create Identity policy, rules, and then define access control policy for controlling access of overlapping users from VRG and VRB.

You have successfully configured user authentication for virtual routers with overlapping networks. VRG and VRB users can now access the shared server through proper NAT translation and authentication controls, with distinct identity policies managing access for each user group.

Interconnect virtual routers using BGP

This task enables you to configure BGP settings to leak routes among virtual routers (Global and user-defined virtual routers) by using route targets and route maps to share routes between different virtual routing instances.

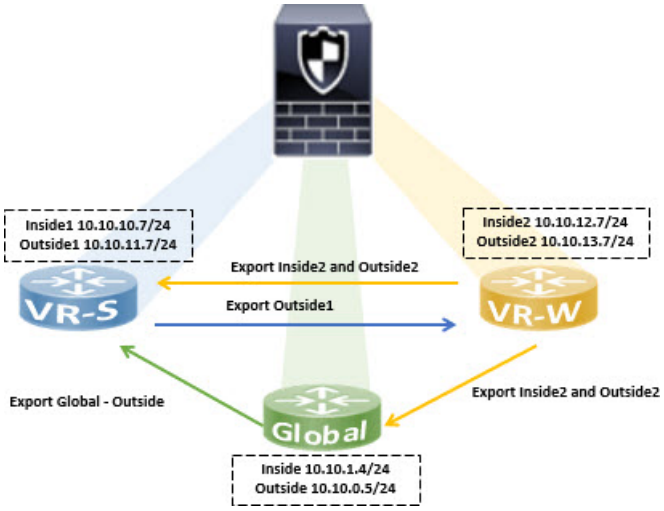
You can now configure BGP settings on a device to leak the routes among virtual routers (Global and user-defined virtual routers). The route target of the source virtual router is exported to the BGP table, which, in turn is imported to the destination virtual router. The route map is used to share the Global virtual routes with the user-defined virtual routers and vice versa. Note that all import or export of the routes to the BGP table is configured at the user-defined virtual router, including the Global virtual routes.

Consider the firewall device of a factory is configured with virtual routers and interfaces:

- Global virtual router is configured with Inside (10.10.1.4/24) and Outside (10.10.0.5/24)
- VR-S (Sales) virtual router is configured with Inside1 (10.10.10.7/24) and Outside1 (10.10.11.7/24)
- VR-W (Warehouse) virtual router is configured with Inside2 (10.10.12.7/24) and Outside2 (10.10.13.7/24)

Assume that you want the routes of warehouse (VR-W) to be leaked with sales (VR-S) and Global, and the outside interface routes of VR-S to VR-W. Similarly, you want the outside interface routes of the Global router to be leaked to sales (VR-S). This example demonstrates the BGP configuration procedure to achieve interconnecting the routers:

Figure 328: Interconnect virtual routers using BGP settings



Before you begin

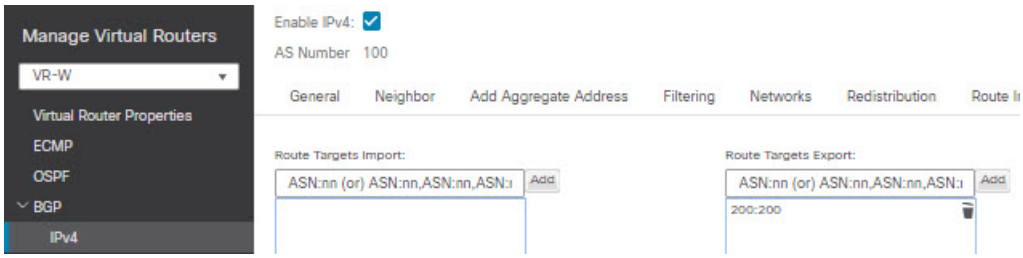
- Create the virtual routers—VR-S and VR-W.
- Enable BGP and for each virtual router Configure BGP for redistribution of connected routes.

Follow these steps to interconnect virtual routers using BGP:

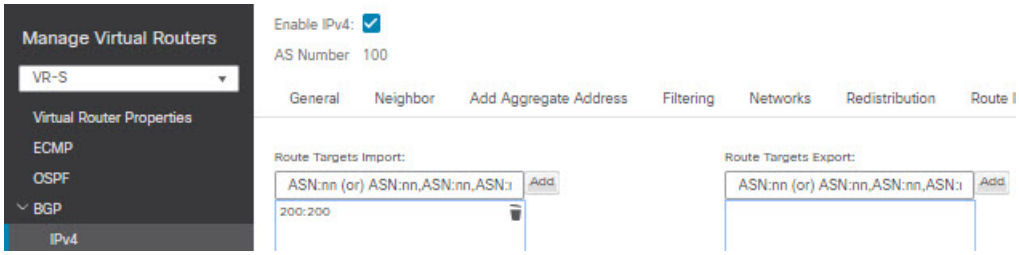
Procedure

Step 1 Configure VR-W to export its routes tagging them with a route target to VR-S:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, edit device, and then click the **Routing** tab.
b) From the virtual router drop-down, select VR-W. Then click **BGP** > **IPv4** > **Route Import/Export**.
c) To leak the VR-W routes to VR-S, tag the routes with a route target, so that the VR-W routes are exported to its BGP table with the route target marked on them. In the **Route Targets Export** field, enter a value, say, *200:200*. Click **Add**:



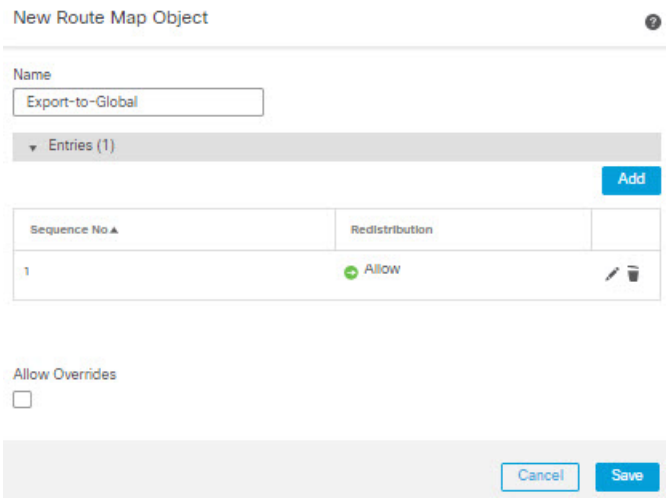
d) From the virtual router drop-down, select VR-S. Then click **BGP** > **IPv4** > **Route Import/Export**.
e) To receive the leaked routes from VR-W, configure the Import Route Target to import the VR-W routes that are marked with the route target from the (peer or redistributed) BGP table. In the **Route Targets Import** field, enter the same route target value that you had configured for VR-W, *200:200*. Click **Add**.



Note
If you want to conditionalize routes to be leaked from VR-W, you can specify the match criteria in the route map object, and choose it in the **User Virtual Router Export Route Map**. Similarly, if you want to conditionalize the routes to be imported to VR-S from the BGP table, you can use the **User Virtual Router Import Route Map**. This procedure is explained in Step 3.

Step 2 Configure VR-W to export its routes to the Global virtual router:

a) You need to create a route map that would allow the VR-W routes to be exported to the Global routing table. In the Cloud-Delivered Firewall Management Center, choose **Objects** > **Object Management** > **Route Map**.
b) Click **Add Route Map**, give a name, say *Export-to-Global*, and then click **Add**.
c) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list:

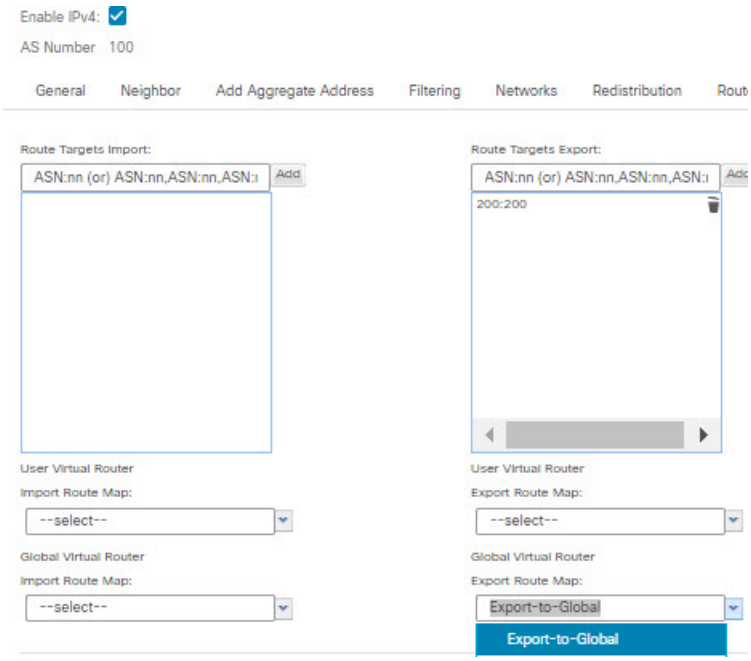


d) Click **Save**.

In this example, all the VR-W routes are leaked to the Global routing table. Hence, no match criteria is configured for the route map.

e) Navigate to the **Routing** tab of the device, and select VR-W. Click **BGP** > **IPv4** > **Route Import/Export**.

f) From the **Global Virtual Router Export Route Map** drop-down list, choose Export-to-Global:



Step 3 To leak only the Outside1 routes of VR-S to VR-W:

a) From the virtual router drop-down, select VR-S and then click **BGP** > **IPv4** > **Route Import/Export**.

b) To leak the VR-S routes to VR-W, tag the routes with a route target, so that the VR-S routes are exported to its BGP table with the route target marked on them. In the **Route Targets Export** field, enter a value, say, *100:100*. Click **Add**.

c) From the virtual router drop-down, select VR-W, and choose **BGP** > **IPv4** > **Route Import/Export**. To receive the leaked routes from VR-S, configure the Import Route Target to import the VR-S routes that are marked with the route target from the (peer or redistributed) BGP table. In the **Route Targets Import** field, enter the VR-S route target value, *100:100*. Click **Add**.

d) Now, you need to conditionalize that only the Outside1 routes of VR-S to be leaked to VR-W. Choose **Objects** > **Object Management** > **Prefix Lists** > **IPv4 Prefix List**. Click **Add IPv4 Prefix List**, give a name, say *VRS-Outside1-Only*, and then click **Add**.

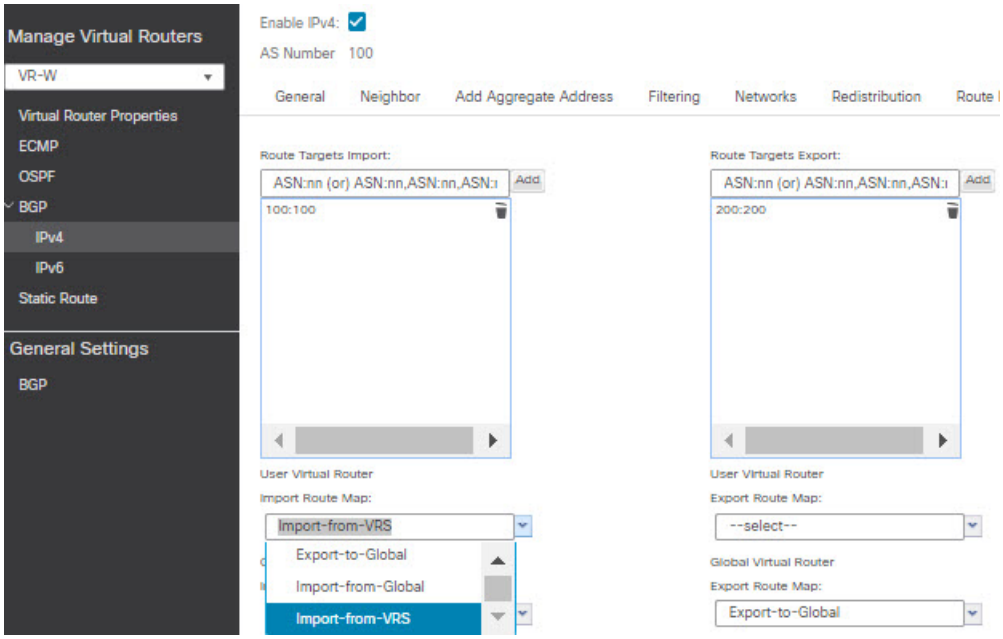
e) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list. Enter the IP Address (first two octets) of the VR-S Outside1 interface and then click **Save**.

f) Create a route map with the match clause with the prefix list. Click **Route Map**. Click **Add Route Map**, give a name, say *Import-from-VRS*, and then click **Add**.

g) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list.

h) In the **Match Clause** tab, click **IPv4**. Under **Address** tab, click **Prefix List**. Under **Available IPv4 Prefix List**, select VRS-Outside1-Only, and then click **Add**.

i) Navigate to the **Routing** tab of the device, and select VR-W. Click **BGP** > **IPv4** > **Route Import/Export**. From the **Global Virtual Router Import Route Map** drop-down list, choose Import-from-VRS:



Step 4 Configure VR-S to import the Outside routes of Global virtual router:

Note
To leak routes to or from a Global virtual router, you must configure the source or destination user defined virtual router respectively. Thus, in this example, VR-S is the destination router that imports the routes from Outside interface of the Global virtual router.

a) In the Cloud-Delivered Firewall Management Center, choose **Objects** > **Object Management** > **Prefix Lists** > **IPv4 Prefix List**.

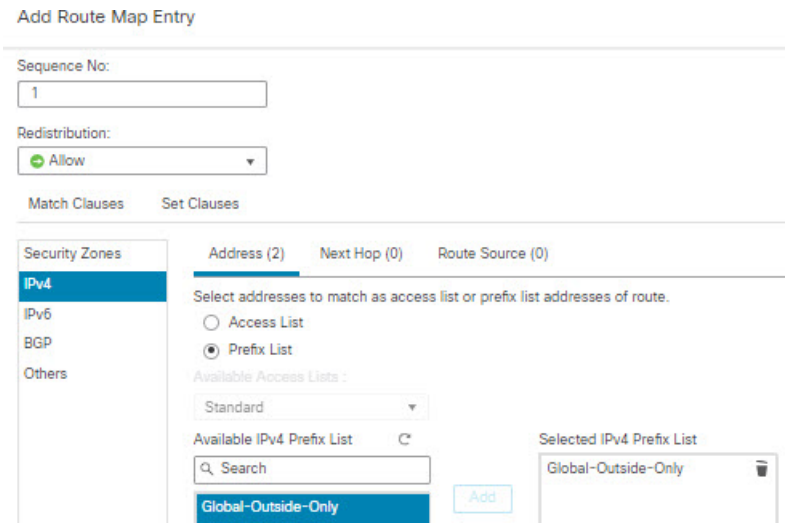
b) Click **Add IPv4 Prefix List**, give a name, say *Global-Outside-Only*, and then click **Add**. Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list.

c) Enter the IP Address (first two octets) of the Global Outside interface:



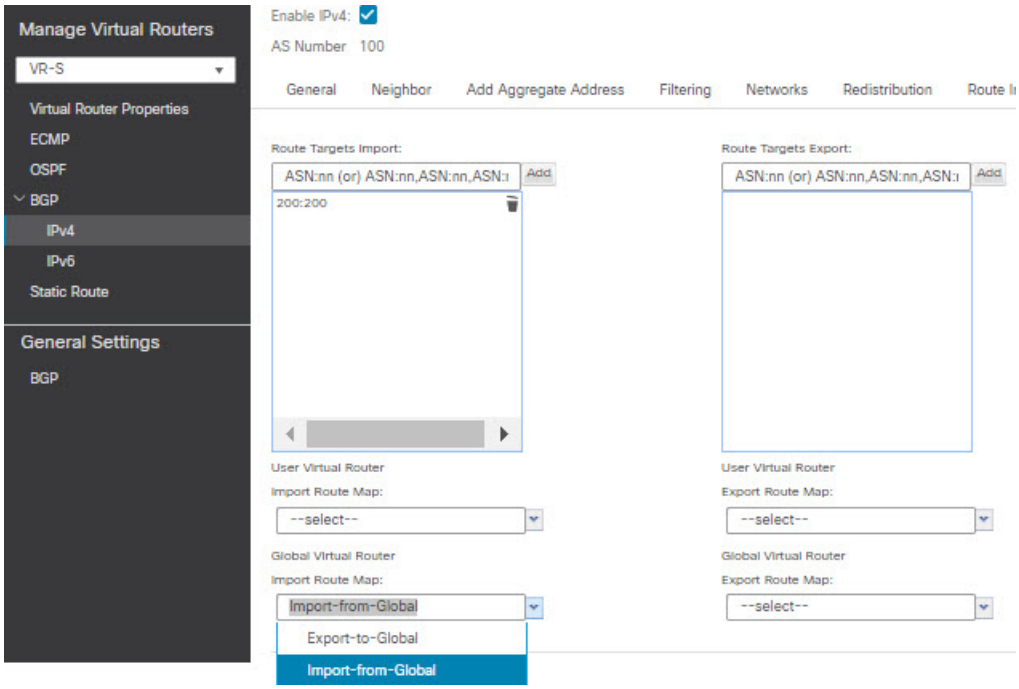
Then click **Save**.

d) Click **Route Map**. Click **Add Route Map**, give a name, say *Import-from-Global*, and then click **Add**.
e) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list.
f) In the **Match Clause** tab, click **IPv4**. Under **Address** tab, click **Prefix List**.
g) Under **Available IPv4 Prefix List**, select Global-Outside-Only, and then click **Add**:



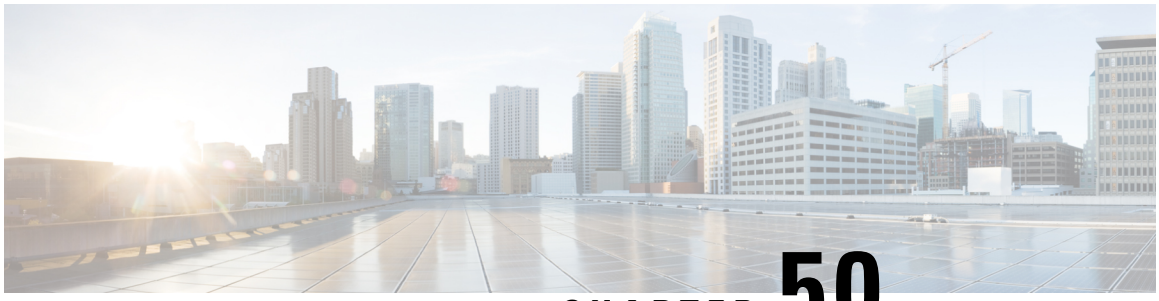
Then click **Save**.

h) Navigate to the **Routing** tab of the device, and select VR-S. Click **BGP** > **IPv4** > **Route Import/Export**.
i) From the **Global Virtual Router Import Route Map** drop-down list, choose Import-from-Global:



Step 5 **Save** and **Deploy**.

The BGP configuration is complete and routes are now leaked among the virtual routers according to the configured route targets and route maps. VR-W routes are exported to VR-S and Global, VR-S Outside1 routes are exported to VR-W, and Global Outside routes are imported to VR-S.



CHAPTER 50

ECMP

Equal Cost Multi-Path (ECMP) is a routing strategy that allows a router to forward packets along multiple paths of equal cost to a destination.

- Enables load balancing across multiple network paths.
- Utilizes routing protocols to determine path cost.
- Increases total available bandwidth for traffic flows.

- ECMP routing methods, on page 1431
- Guidelines and limitations for ECMP, on page 1432
- Manage ECMP zones for virtual routers, on page 1433
- Create an ECMP zone, on page 1433
- Configure an equal cost static route, on page 1434
- Modify an ECMP zone, on page 1435
- Remove an ECMP zone, on page 1436
- Configure ECMP zones on your device, on page 1436

ECMP routing methods

ECMP is a routing method that

- enables traffic distribution across multiple equal-cost paths,
- supports up to 8 equal cost static or dynamic routes across up to 8 interfaces within each zone, and
- uses traffic zones per virtual router to contain a group of interfaces.

Multiple default routes configuration

This example shows multiple default routes configured across three interfaces in the zone:

```
route for 0.0.0.0 0.0.0.0 through outside1 to 10.1.1.2

route for 0.0.0.0 0.0.0.0 through outside2 to 10.2.1.2

route for 0.0.0.0 0.0.0.0 through outside3 to 10.3.1.2
```

Guidelines and limitations for ECMP

Follow these principles for effective ECMP zone configuration:

Firewall mode limitations

Use ECMP zones only in routed firewall mode.

Interface restrictions

Do not use dVTI or Loopback interfaces with ECMP zones.

ECMP zone configuration limits

Follow these limits when configuring ECMP zones:

- Devices can have a maximum of 256 ECMP zones.
- You can associate only 8 interfaces per ECMP zone.
- An interface can be a member of only one ECMP zone.

Interface management limitations

Do not remove interfaces or delete zones actively used for routing:

- You cannot remove an interface that is associated with an equal cost static route from the ECMP zone.
- You cannot delete an ECMP zone if its interface has equal cost static routes associated with it.

Supported interface types

Use only routed interfaces for ECMP zones. Do not associate these interface types with an ECMP zone:

- BVI interface.
- Member interfaces in an EtherChannel.
- Failover or state link interface.
- Management-only or management-access interfaces.
- Cluster control link interface.
- VNIs.
- VLAN interfaces.
- Interfaces in a remote access VPN configuration with SSL enabled.

Feature compatibility limitations

Consider these feature limitations when using ECMP zones:

- DHCP Relay is not supported on interfaces in an ECMP zone.
- Dual ISP/WAN Firewall Threat Defense Deployment—Create a single ECMP zone for the primary and secondary data interfaces, enabling static routes with identical metric values.

- The Firewall Threat Defense does not support ECMP with NAT in IPsec sessions—a standard IPsec virtual private network (VPN) tunnel does not work with NAT points in the delivery path of IPsec packets.

Manage ECMP zones for virtual routers

The ECMP page lets you manage existing ECMP zones linked with virtual routers. When you click **ECMP** on the **Routing** pane, the ECMP page displays the current ECMP zones, including the associated interfaces, for the selected virtual router. On this page, you can add new ECMP zones to the virtual router. You can also **Edit** ( ) and **Delete** ( ) ECMP zones.

You can perform these actions:

- Create an ECMP zone, on page 1433
- Configure an equal cost static route, on page 1434
- Modify an ECMP zone, on page 1435
- Remove an ECMP zone, on page 1436

Create an ECMP zone

Create an ECMP zone to enable load balancing across multiple equal-cost paths in a virtual router configuration.

ECMP zones are created per virtual router. Only interfaces that belong to the virtual router where the ECMP is being created can be associated with the ECMP zone.

Procedure

Step 1 Navigate to **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 From the virtual router drop-down, select the virtual router in which you want to create the ECMP zone.

You can create ECMP zones in the global virtual router and user-defined virtual routers. For information on creating virtual routers, refer to Create a virtual router, on page 1386.

Step 4 Click **ECMP**.

Step 5 Click **Add**.

Step 6 In the **Add ECMP** box, enter a name for the ECMP zone.

Note
The ECMP name must be unique on the routed device.

Step 7 Under the **Available Interfaces** box, select the interface and click **Add**.

- Only interfaces belonging to the virtual router where you are creating the ECMP can be associated with it. Interfaces must have a logical name to be displayed; edit the interface and save to set a logical name.

- From Version 10.0, ECMP is supported on dynamic VTIs (DVTIs) of hub devices in SD-WAN VPN topologies. When enabled on a DVTI, it is automatically added to a system-generated ECMP zone, so DVTIs do not appear under **Available Interfaces**.

Step 8 Click **OK**.

The ECMP page displays your newly created ECMP zone.

Step 9 Click **Save** and **Deploy** the configuration.

You can associate the ECMP zone interfaces with equal cost static route by defining them with same destination and metric value but with a different gateway.

What to do next

- Configure an equal cost static route, on page 1434
- Modify an ECMP zone, on page 1435
- Remove an ECMP zone, on page 1436

Configure an equal cost static route

Configure equal cost static routes to enable load balancing across multiple interfaces with the same destination and metric value within an ECMP zone.

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	N/A	Firewall Threat Defense and Firewall Threat Defense Virtual	Any	Admin/Network Admin/Security Approver

You can assign interfaces of a virtual router, both global and user-defined, to an ECMP zone for the device.

Before you begin

- To configure an equal cost static route for an interface, ensure that it is associated with an ECMP zone. For more information, refer to Create an ECMP zone, on page 1433.
- All routing configuration settings of a non-VRF-capable device are also available for a global virtual router.
- You cannot define a static route for interfaces with the same destination and metric without associating the interfaces with an ECMP zone.

Procedure

Step 1 From the **Devices** > **Device Management** page, edit the Firewall Threat Defense device. Click the **Routing** tab.

Step 2 From the drop-down list, select the virtual router whose interfaces are associated with an ECMP zone.

Step 3 To configure the equal cost static route for the interfaces, click **Static Route**.

Step 4 Either click **Add Route** to add a new route, or click **Edit** (✎) for an existing route.

Step 5 From the **Interface** drop-down, select the interface that belongs to the virtual router and is associated with an ECMP zone.

Step 6 Select the destination network from the **Available Networks** box and click **Add**.

Step 7 Enter a gateway for the network.

Step 8 Enter a metric value. It can be a number that ranges between 1 and 254.

Step 9 To save the settings, click **Save**.

To configure equal cost static routing, repeat these steps to configure the static route for another interface in the same ECMP zone with the same destination network and metric value. Remember to provide a different gateway.

The equal cost static routes are configured for the interfaces associated with the ECMP zone, enabling load balancing across multiple paths.

What to do next

- Modify an ECMP zone, on page 1435
- Remove an ECMP zone, on page 1436

Modify an ECMP zone

Use this procedure when you need to make changes to an existing ECMP zone configuration on your device.

Follow these steps to modify an ECMP zone:

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **ECMP**.

ECMP zones with their associated interfaces appear.

Step 4 To modify an ECMP zone, click **Edit** (✎) next to the desired ECMP zone. In the **Edit ECMP** dialog box, you can make changes to the ECMP name or associated interfaces.

- ECMP Name: Ensure that the ECMP name you enter is unique.
- Interfaces: Add or remove one or more interfaces as needed.

Note
An interface must not belong to more than one ECMP zone, and it should not be linked to an equal-cost static route.

Step 5 Click **OK**.

Step 6 Click **Save**.

The ECMP zone shows your updated settings.

What to do next

- Configure an equal cost static route, on page 1434
- Remove an ECMP zone, on page 1436

Remove an ECMP zone

ECMP zones group interfaces for equal-cost multipath routing. You may need to remove an ECMP zone when reconfiguring your network topology or when the zone is no longer required.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **ECMP**.

The ECMP zones with associated interfaces appear.

Step 4 To remove an ECMP zone, click **Delete** ( ) next to the ECMP zone.

Note
You cannot delete the ECMP zone if it has interfaces associated with equal-cost static routes.

Step 5 Click **Delete** in the confirmation message.

Step 6 Click **Save** to apply changes.

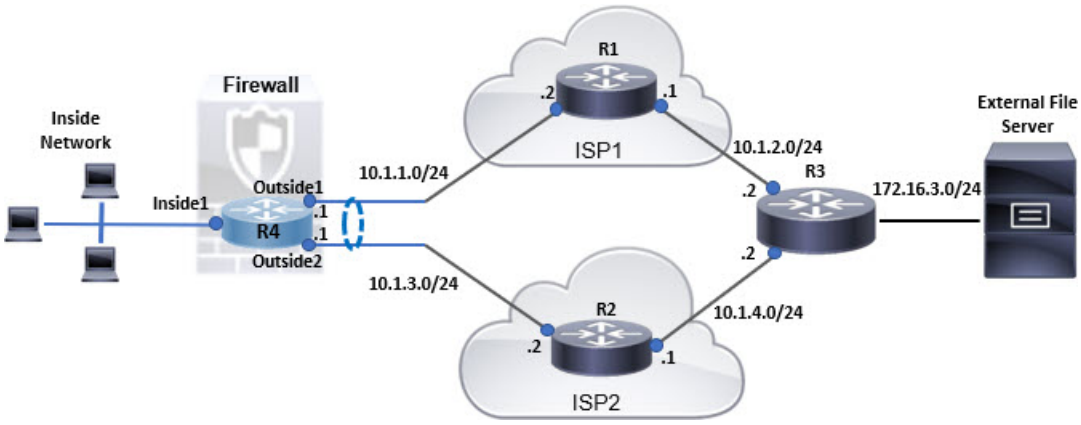
The ECMP zone is removed from the device configuration and is no longer available for routing operations.

Configure ECMP zones on your device

Configure Equal-Cost MultiPath (ECMP) to enable efficient traffic handling through the device with support for asymmetric routing, load balancing, and traffic rerouting when a path becomes unavailable.

This example demonstrates how to use Cloud-Delivered Firewall Management Center to configure ECMP zones on Firewall Threat Defense such that the traffic flowing through the device is handled efficiently. With ECMP configured, Firewall Threat Defense maintains the routing table on a per-zone basis, enabling efficient packet re-routing. Thus, ECMP supports asymmetric routing, load balancing, and traffic rerouting when a path becomes unavailable. In this example, R4 records the two paths to reach the external file server.

Figure 329: Configuration example for ECMP



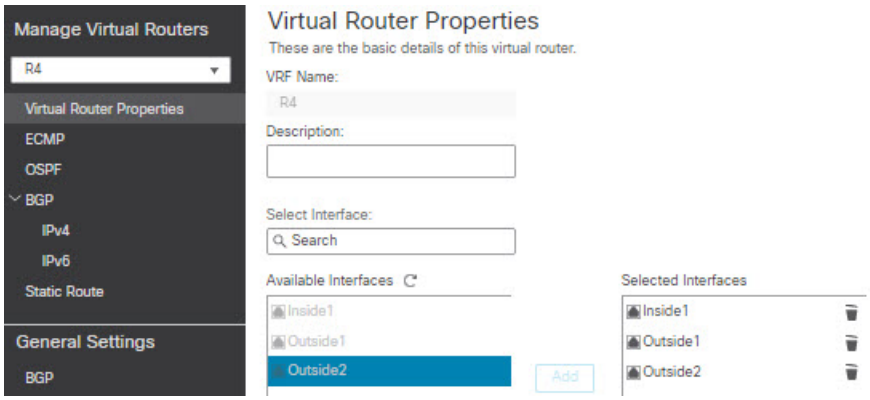
Follow these steps to configure ECMP on your device:

Procedure

Step 1 Create a virtual router.

Set up a new router on *R4* with interfaces: *Inside1*, *Outside1*, and *Outside2*. For more information, refer to Create virtual router.

Figure 330: Configuring R4 virtual router



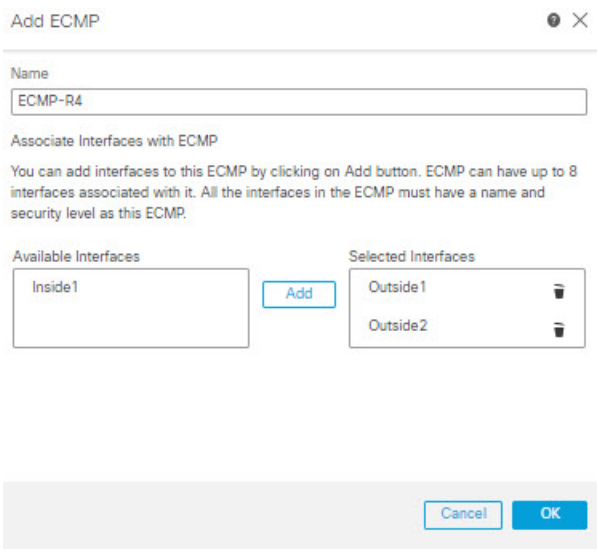
Step 2 Create ECMP zones:

a) In the **Routing** tab, choose R4 user-defined virtual router, and then click **ECMP**.

b) Click **Add**.

c) Enter an ECMP name and from the **Available Interfaces** list, choose *Outside1* and *Outside2*:

Figure 331: Creating ECMP zone

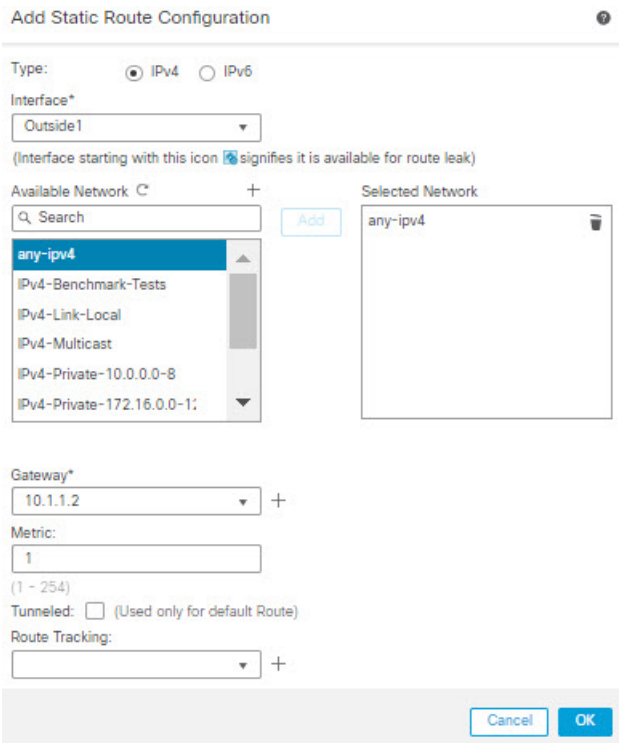


d) Click **OK**.
e) Click **Save**.

Step 3 Create static routes for the zone interfaces:

a) In the **Routing** tab, click **Static Route**.
b) From the **Interface** drop-down list, select Outside1.
c) Under **Available Network**, choose any-ipv4 and click **Add**.
d) Specify the next-hop address in the **Gateway** field, 10.1.1.2.

Figure 332: Configuring static route for Outside1



e) Configure the static route for Outside2 by repeating Step 3b through Step 3d.

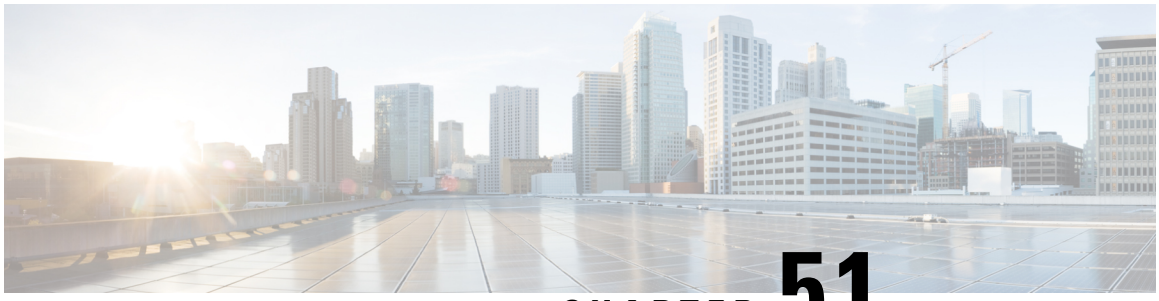
Ensure that you specify the same metric but different gateways for the static routes:

Figure 333: Configured static routes of ECMP zone interfaces



Step 4 Save and deploy the changes to apply ECMP routing to your device.

Network packets will now utilize efficient routes to reach their destination R3, either through R4>R1>R3 or R4>R2>R3, according to the ECMP algorithm configuration. If the R1>R3 route becomes unavailable, traffic can be rerouted through R2, which helps reduce disruption during path failure. Additionally, return traffic from R3 can arrive on *Outside2* even if the original packet was sent through *Outside1*. When network traffic is heavy, R4 distributes the network load between two specified routes to maintain balanced traffic.



CHAPTER 51

Bidirectional Forwarding Detection Routing

This chapter describes how to configure Firewall Threat Defense to use the Bidirectional Forwarding Detection (BFD) routing protocol.

- BFD routing protocol, on page 1441
- Guidelines for BFD routing, on page 1442
- Configure BFD, on page 1443
- History for BFD routing, on page 1446

BFD routing protocol

A BFD routing protocol is a detection protocol that

- provides fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols
- operates in a unicast, point-to-point mode on top of any data protocol being forwarded between two systems, and
- carries packets in the payload of the encapsulating protocol appropriate for the media and the network.

Why BFD?

BFD provides a consistent failure detection method for network administrators in addition to fast forwarding path failure detection.

Network administrators can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms. This makes network profiling and planning easier, and reconvergence time is consistent and predictable.

Note In Firewall Threat Defense, BFD is supported on BGP protocols only.

Guidelines for BFD routing

Context mode guidelines

BFD is supported on all Firewall Threat Defense platforms and is available in multi-instance mode.

Firewall mode guidelines

BFD is supported in routed firewall mode and not in transparent mode.

Failover and cluster guidelines

- BFD is not supported on failover interfaces.
- In clustering, BFD is supported only on the control node.

Routing and protocol guidelines

- OSPFv2, OSPFv3, BGP IPv4, and BGP IPv6 protocol are supported.

  IS-IS and EIGRP protocols are not supported.
- BFD for static routes is not supported. You can configure BFD on interfaces that belong only to virtual routers.
- Only named interfaces are supported.
- BFD on BVI, VTI, and loopback interfaces are not supported.

Single-hop guidelines

- Echo mode is disabled by default. You can enable echo mode on single-hop only.
- Echo mode is not supported for IPv6.
- Use only a single-hop template to configure a single-hop policy.
- Authentication of the single-hop template is optional.
- You cannot configure multiple BFDs on the same interface.

Multi-hop guidelines

- Do not configure the source IP address also as the destination IP address.
- Source and destination address should have same IP type—IPV4 or IPV6.
- Only network objects of host or network type are allowed.
- Use only a multi-hop template to configure a multi-hop policy.
- Authentication is mandatory for the multi-hop template.

Upgrade guidelines

When you upgrade to version 7.3 and the previous version has any FlexConfig BFD policies, the management center displays a warning message during deployment. However, it does not stop the deployment process. After post-upgrade deployment, to manage the BFD policies from the UI (**Device (Edit)** > **Routing** > **BFD**), you must configure BFD policies in the **Device (Edit)** > **Routing** > **BFD** page and remove the configuration from the FlexConfig policy for the device.

Configure BFD

This task enables and configures the Bidirectional Forwarding Detection (BFD) routing policy on your system to provide rapid detection of faults in the forwarding path.

Before you begin

Ensure you have access to the system where you want to configure BFD.

Follow these steps to configure BFD:

Procedure

Step 1 Create Configure a BFD template, on page 2956.

Step 2 Configure BFD policies, on page 1443.

Step 3 Configure BFD support in the BGP neighbor settings; see, Configure BGP Neighbor Settings.

BFD is enabled and configured on your system.

Configure BFD policies

BFD policies are used to provide rapid failure detection for interfaces associated with virtual routers. Configuring BFD templates allows you to monitor connectivity and quickly identify link failures.

- Enable fast failure detection for routing interfaces.
- Support both single-hop and multi-hop BFD configurations.

You can bind a BFD template to an interface belonging to a virtual router, or to a source and destination address pair.

Before you begin

- BFD policy is configurable only on interfaces that belong to virtual routers. See Configure interfaces to a virtual router.

To configure BFD policies, complete these steps.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, from the **Devices** > **Device Management** page, edit the virtual-router supported device. Navigate to **Routing**.

Step 2 From the drop-down list, select the desired virtual router, and then click **BFD**.

Step 3 To configure a BFD on the interface, click the **Single-Hop** tab or **Multi-Hop** tab.

Note
For a single-hop policy, the BFD template is configured on an interface; for a multi-hop policy, the BFD template is configured on a source and destination address pair.

Step 4 Click **Add**. To modify the configured BFD policy, click **Edit** ( ).

Note
When you edit the interface mapping with BFD template to replace it with a new BFD template, the Cloud-Delivered Firewall Management Center uses a **no** command to remove the template mapping from interface and applies the new template to the interface which causes a BFD flap which may also lead to an OSPFv2, OSPFv3, or BGP flap. However, if the BFD intervals are higher, the BFD flap might not occur. Alternatively, to avoid the flapping, you can delete the existing BFD template mapping; deploy the interface, and then add the new BFD template to the interface and deploy the configuration.

- See Configure a single-hop BFD policy, on page 1444.
- See Configure a multi-hop BFD policy, on page 1445.

After completing this task, BFD policies are configured on the selected interfaces or address pairs.

Configure a single-hop BFD policy

You can configure a single-hop BFD policy only on an interface that belongs to a virtual router. This task is relevant when you need to apply BFD for rapid link failure detection on specific interfaces.

Before you begin

- Create single-hop BFD templates. You cannot configure single-hop BFD policy on interfaces that use a multi-hop template.

Follow these steps to configure a single-hop BFD policy:

Procedure

Step 1 In the **Single-Hop** tab, click **Add** or **Edit**.

Step 2 In the **Add BFD Single-Hop** dialog box, configure the following:

a) The **Interface** drop-down list displays interfaces belonging to virtual routers. Select the interface you want to configure with the BFD policy.

b) The **Template Name** drop-down list displays single-hop templates. Select the template that you want to apply.

If you have not created a single-hop template, use **Add** (+) and create a single-hop BFD template.

Step 3 Click **OK** and **Save** the configuration.

The single-hop BFD policy is applied to the selected interface. The interface will now use the specified BFD template for single-hop operations.

Configure a multi-hop BFD policy

Configure a multi-hop BFD policy to enable monitoring of connectivity between specific source and destination address pairs. This configuration helps to monitor links that traverse multiple hops in your network.

Before you begin

- Create multi-hop BFD templates. You cannot configure multi-hop BFD policy using a single-hop template.

Follow these steps to configure a multi-hop BFD policy:

Procedure

Step 1 In the **Add BFD Multi-Hop** dialog box, configure the following:

a) Click the BFD source address type:**IPv4** or **IPv6** radio button.

b) The **Source Address** drop-down list displays the network objects. Select the source address for the BFD policy; do not choose *any-ipv4* or *any-ipv6*.

If you have not created the required network object, use **Add** (+) and create a host/network object.

Note
The created network object's IP type should match with the selected source IP type.

c) The **Destination Address** drop-down list displays network objects. Select the destination address for the BFD; do not choose *any-ipv4* or *any-ipv6*.

If you have not created the required network object, use **Add** (+) and create a host/network object.

Note
The created network object's IP type should match with the selected source IP type.

Attention
Do not select the network object that has the same IP address as that of the source address.

d) The **Template Name** drop-down list displays multi-hop templates. Select the template to apply to the BFD policy.

If you have not created a multi-hop template, use **Add** (+) and create a multi-hop BFD template.

Step 2 Click **OK** and **Save** the configuration.

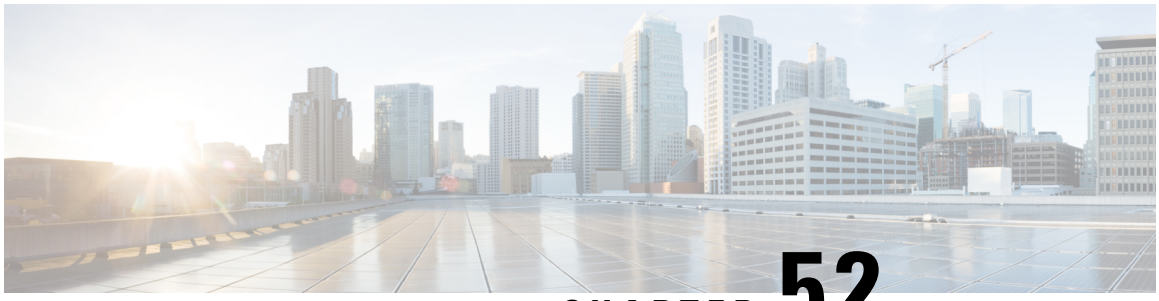
The multi-hop map (table view) is displayed on the **Multi-Hop** tab page.

History for BFD routing

The History table summarizes the introduction and evolution of the BFD routing in Cisco Cloud-Delivered Firewall Management Center across releases, enabling quick tracking of feature availability and changes.

Table 93: Feature history for BFD routing

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
BFD support for OSPF	7.4	7.4	You can enable BFD on interfaces that use OSPF version 2 and OSPF version 3. New/modified screens: • **Configuration** > **Device Setup** > **Routing** > **OSPFv2** • **Configuration** > **Device Setup** > **Routing** > **OSPFv3**
BFD configuration	7.4	7.4	In the previous releases, BFD was configurable on threat defense only through FlexConfig. FlexConfig no longer supports BFD configuration. You can now configure BFD policies for threat defense in the management center UI. In threat defense, BFD is supported only on the BGP protocol. New/modified screens: **Devices** > **Device Management** > **Routing** > **BFD**.



CHAPTER 52

OSPF

This chapter describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

- OSPF, on page 1447
- Requirements and prerequisites for OSPF, on page 1451
- Guidelines for OSPF, on page 1451
- OSPFv2 configuration, on page 1454
- OSPFv3 configuration, on page 1467
- History for OSPF, on page 1477

OSPF

This chapter describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

OSPF routing

OSPF is an interior gateway routing protocol that

- uses link states rather than distance vectors for path selection,
- propagates link-state advertisements rather than routing table updates, and
- converges more quickly than RIP networks because only LSAs are exchanged instead of the entire routing tables.

OSPF characteristics and features

OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area contains an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

The advantages of OSPF over RIP include these characteristics:

- OSPF link-state database updates are sent less frequently than RIP updates, and the link-state database is updated instantly, rather than gradually, AS stale information is timed out.

- Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The Firewall Threat Defense device calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.

The disadvantage of shortest path first algorithms is that they require a lot of CPU cycles and memory.

The Firewall Threat Defense device can run two processes of OSPF protocol simultaneously on different sets of interfaces. You might want to run two processes if you have interfaces that use the same IP addresses (NAT allows these interfaces to coexist, but OSPF does not allow overlapping addresses). Or you might want to run one process on the inside and another on the outside, and redistribute a subset of routes between the two processes. Similarly, you might need to segregate private addresses from public addresses.

You can redistribute routes into an OSPF routing process from another OSPF routing process, a RIP routing process, or from static and connected routes configured on OSPF-enabled interfaces.

The Firewall Threat Defense device supports these OSPF features:

- Intra-area, inter-area, and external (Type I and Type II) routes.
- Virtual links.
- LSA flooding.
- Authentication to OSPF packets (both password and MD5 authentication).
- Configuring the Firewall Threat Defense device AS a designated router or a designated backup router. The Firewall Threat Defense device also can be set up AS an ABR.
- Stub areas and not-so-stubby areas.
- Area boundary router Type 3 LSA filtering.

OSPF supports both MD5 and clear text neighbor authentication. Authentication should be used with all routing protocols when possible because route redistribution between OSPF and other protocols (such AS RIP) can potentially be used by attackers to subvert routing information.

If NAT is used, if OSPF is operating on public and private areas, and if address filtering is required, then you need to run two OSPF processes—one process for the public areas and one for the private areas.

A router that has interfaces in multiple areas is called an Area Border Router (ABR). A router that acts AS a gateway to redistribute traffic between routers using OSPF and routers using other routing protocols is called an Autonomous System Boundary Router (ASBR).

An ABR uses LSAs to send information about available routes to other OSPF routers. Using ABR Type 3 LSA filtering, you can have separate private and public areas with the ASA acting AS an ABR. Type 3 LSAs (inter-area routes) can be filtered from one area to other, which allows you to use NAT and OSPF together without advertising private networks.



Note Only Type 3 LSAs can be filtered. If you configure the Firewall Threat Defense device AS an ASBR in a private network, it will send Type 5 LSAs describing private networks, which will get flooded to the entire AS, including public areas.

If NAT is employed but OSPF is only running in public areas, then routes to public networks can be redistributed inside the private network, either AS default or Type 5 AS external LSAs. However, you need to configure

static routes for the private networks protected by the Firewall Threat Defense device. Also, you should not mix public and private networks on the same Firewall Threat Defense device interface.

You can have two OSPF routing processes, one RIP routing process, and one EIGRP routing process running on the Firewall Threat Defense device at the same time.

OSPF support for fast hello packets

OSPF support for fast hello packets is a feature that

- provides a way to configure the sending of hello packets in intervals less than one second, and
- results in faster convergence in an Open Shortest Path First (OSPF) network.

OSPF fast hello packets configuration prerequisites

OSPF must be configured in the network already or configured at the same time as the OSPF Support for Fast Hello Packets feature.

OSPF hello interval and dead interval

OSPF hello interval and dead interval are configurable time parameters that

- maintain connectivity between OSPF neighbors through regular hello packet exchanges,
- detect neighbor failures when hello packets are not received within the specified timeframes, and
- must be consistent across all routers within the same network to ensure proper operation.

Hello interval and dead interval configuration details

OSPF hello packets are packets that an OSPF process sends to its OSPF neighbors to maintain connectivity with those neighbors. The hello packets are sent at a configurable interval (in seconds). The defaults are 10 seconds for an Ethernet link and 30 seconds for a non broadcast link.

Hello packets include a list of all neighbors for which a hello packet has been received within the dead interval. The dead interval is also a configurable interval (in seconds), and defaults to four times the value of the hello interval. The value of all hello intervals must be the same within a network. Likewise, the value of all dead intervals must be the same within a network.

These two intervals work together to maintain connectivity by indicating that the link is operational. If a router does not receive a hello packet from a neighbor within the dead interval, it will declare that neighbor to be down.

OSPF fast hello packets

OSPF fast hello packets are hello packets that

- are sent at intervals of less than 1 second,
- are achieved by using the OSPF dead-interval command with a dead interval set to 1 second, and
- use a hello-multiplier value set to the number of hello packets sent during that 1 second.

Fast hello packet behavior

When fast hello packets are configured on the interface, these behaviors occur:

- The hello interval advertised in the hello packets that are sent out this interface is set to 0.
- The hello interval in the hello packets received over this interface is ignored.

The dead interval must be consistent on a segment, whether it is set to 1 second (for fast hello packets) or set to any other value. The hello multiplier need not be the same for the entire segment as long as at least one hello packet is sent within the dead interval.

To understand fast hello packets, you should already understand the relationship between OSPF hello packets and the dead interval. See OSPF hello interval and dead interval, on page 1449.

Benefits of OSPF fast hello packets

A fast hello packet is an OSPF protocol feature that

- reduces network convergence time by enabling routers to detect changes in the network faster,
- allows rapid detection of lost neighbors within one second, and
- enhances reliability in LAN environments where neighbor loss may not be detected at the Open Systems Interconnection (OSI) physical or data-link layer.

The benefit of OSPF Fast Hello Packets is that your OSPF network can react swiftly to changes, minimizing downtime and improving overall network stability. This feature is particularly useful in scenarios where routers may not detect lost neighbors via traditional methods, allowing OSPF to maintain accurate routing tables.

If a router on a LAN segment is powered off or disconnected and the physical and data-link layers do not signal neighbor loss, OSPF fast hello packets will detect the neighbor's disappearance within one second.

OSPFv2 and OSPFv3 implementation differences

A protocol distinction is a protocol category that

- prevents backward compatibility between OSPFv2 and OSPFv3,
- requires both OSPF versions to run simultaneously for IPv4 and IPv6 routing, and
- allows OSPFv2 and OSPFv3 to coexist without interacting with each other.

OSPFv3 additional features

OSPFv3 provides additional features that include:

- Protocol processing per link.
- Removal of addressing semantics.
- Addition of flooding scope.
- Support for multiple instances per link.
- Use of the IPv6 link-local address for neighbor discovery and other features.

- LSAs expressed as prefix and prefix length.
- Addition of two LSA types.
- Handling of unknown LSA types.
- Authentication support using the IPsec ESP standard for OSPFv3 routing protocol traffic, as specified by RFC-4552.

Requirements and prerequisites for OSPF

This reference provides the requirements and prerequisites that must be met before configuring OSPF on supported devices.

Model support

Firewall Threat Defense

Firewall Threat Defense Virtual

Supported domains

Any

User roles

Admin

Network Admin

Guidelines for OSPF

Context mode guidelines

OSPFv2 supports single and multiple context mode.

- OSPFv2 instances cannot form adjacencies with each other across shared interfaces because, by default, inter-context exchange of multicast traffic is not supported across shared interfaces. However, you can use the static neighbor configuration under OSPFv2 process configuration under OSPFv2 process to bring up OSPFv2 neighbourship on a shared interface.
- Inter-context OSPFv2 on separate interfaces is supported.
- **(For point-to-point topology)** In a multiple context mode, you can configure static OSPFv2 neighbors to advertise OSPFv2 routes over a point-to-point, non-broadcast network. For a successful OSPF adjacency formation, when you remove and reconfigure OSPF on a shared interface, execute the commands in this order: Remove the OSPF configuration using **no router <OSPF name>**, and then remove the OSPF point-to-point non-broadcast configuration on the interface using **no OSPF network point-to-point non-broadcast**. When you reconfigure OSPF on the shared interface, configure the OSPF router using **router <OSPF name>**, and then configure the interface with **OSPF network point-to-point non-broadcast**.
- OSPFv3 supports single mode only.

Key chain authentication guidelines

OSPFv2 supports key chain authentication on both single and multiple mode, both in physical and virtual modes. However, in multiple mode, you can configure the key chain only in context mode.

- The rotating keys are applicable only for OSPFv2 protocol. OSPF area authentication with key chain is not supported.
- The existing MD5 authentication without time range in OSPFv2 is still supported along with new rotating keys.
- Though the platform supports SHA1 and MD5 cryptographic algorithms, only MD5 cryptographic algorithm is used for authentication.

Firewall mode guidelines

OSPF supports routed firewall mode only. OSPF does not support transparent firewall mode.

High availability Guidelines

OSPFv2 and OSPFv3 support Stateful High availability.

IPv6 guidelines

OSPFv2 does not support IPv6, while OSPFv3 provides IPv6 support with specific authentication and routing capabilities.

- OSPFv2 does not support IPv6.
- OSPFv3 supports IPv6.
- OSPFv3 uses IPv6 for authentication.
- The Firewall Threat Defense device installs OSPFv3 routes into the IPv6 RIB, provided it is the best route.

OSPFv3 hello packets and GRE

Typically, OSPF traffic does not pass through GRE tunnel. When OSPFv3 on IPv6 is encapsulated inside GRE, the IPv6 header validation for security check such as Multicast Destination fails. The packet is dropped due to the implicit security check validation, as this packet has destination IPv6 multicast.

- You may define a pre-filter rule to bypass GRE traffic. However, with pre-filter rule, inner packets would not be interrogated by the inspection engine.

Clustering guidelines

OSPFv2 and OSPFv3 have specific limitations and behaviors in clustering environments.

- OSPFv3 encryption is not supported. An error message appears if you try to configure OSPFv3 encryption in a clustering environment.
- In Spanned interface mode, dynamic routing is not supported on management-only interfaces.
- In Individual interface mode, make sure that you establish the control and data units as either OSPFv2 or OSPFv3 neighbors.

- In Individual interface mode, OSPFv2 adjacencies can only be established between two contexts on a shared interface on the control unit. Configuring static neighbors is supported only on point-to point-links; therefore, only one neighbor statement is allowed on an interface.
- When a control role change occurs in the cluster, specific behavior occurs based on interface mode. In spanned interface mode, the router process is active only on the control unit and is in a suspended state on the data units. Each cluster unit has the same router ID because the configuration has been synchronized from the control unit. As a result, a neighboring router does not notice any change in the router ID of the cluster during a role change. In individual interface mode, the router process is active on all the individual cluster units. Each cluster unit chooses its own distinct router ID from the configured cluster pool. A control role change in the cluster does not change the routing topology in any way.

Multiprotocol label switching (MPLS) and OSPF guidelines

When a MPLS-configured router sends Link State (LS) update packets containing opaque Type-10 link-state advertisements (LSAs) that include an MPLS header, authentication fails and the appliance silently drops the update packets, rather than acknowledging them. Eventually the peer router will terminate the neighbor relationship because it has not received any acknowledgments.

- Make sure that non-stop forwarding (NSF) is disabled on the appliance to ensure that the neighbor relationship remains stable. Navigate to the **Non Stop Forwarding** page in Cloud-Delivered Firewall Management Center (**Devices** > **Device Management (select the desired device)** > **Routing** > **OSPF** > **Advanced** > **Non Stop Forwarding**). Ensure the **Non Stop Forwarding Capability** boxes are not checked.

These guidelines apply to MPLS environments where opaque LSAs might cause authentication failures.

Note The Firepower 4100/9300 models may have high latency when using MPLS because they lack load balancing across multiple receiving queues.

Bidirectional and forwarding detection (BFD) and OSPF guidelines

You can enable BFD on specific interface types with OSPFv2 and OSPFv3.

- You can enable BFD on OSPFv2 and OSPFv3 interfaces (Physical Interfaces, Sub-Interfaces, and Port-Channels).
- BFD is not supported on VTI Tunnels, DVTI Tunnels, Loopback, Switchport, VNI, VTEP, and IRB interfaces.

Route redistribution guidelines

Route redistribution with prefix lists and between routing domains requires specific configuration considerations.

- Redistribution of route maps with IPv4 prefix list on OSPFv2 is supported. However, redistribution of route maps with IPv6 prefix list on OSPFv3 is not supported. Use an access list in the route map on OSPF for redistribution.
- When OSPF is configured on a device that is a part of EIGRP network or vice versa, ensure that OSPF-router is configured to tag the route (EIGRP does not support route tag yet). When redistributing OSPF into EIGRP and EIGRP into OSPF, a routing loop occurs when there is an outage on one of the links, interfaces, or even when the route originator is down. To prevent the redistribution of routes from

one domain back into the same domain, a router can tag a route that belongs to a domain while it is redistributing, and those routes can be filtered on the remote router based on the same tag. Because the routes will not be installed into the routing table, they will not be redistributed back into the same domain.

Additional guidelines

OSPFv2 and OSPFv3 support various features and have specific limitations in different deployment scenarios.

- OSPFv2 and OSPFv3 support multiple instances on an interface.
- OSPFv3 supports encryption through ESP headers in a non-clustered environment.
- OSPFv3 supports Non-Payload Encryption.
- OSPFv2 supports Cisco NSF Graceful Restart and IETF NSF Graceful Restart mechanisms as defined in RFCs 4811, 4812 & 3623 respectively.
- OSPFv3 supports Graceful Restart mechanism as defined in RFC 5187.
- There is a limit to the number of intra area (type 1) routes that can be distributed. For these routes, a single type-1 LSA contains all prefixes. Because the system has a limit of 35 KB for packet size, 3000 routes result in a packet that exceeds the limit. Consider 2900 type 1 routes to be the maximum number supported.
- For a device using virtual routing, you can configure OSPFv2 and OSPFv3 for a global virtual router. However, you can configure only OSPFv2 for a user-defined virtual router.
- To avoid adjacency flaps due to route updates being dropped if the route update is larger than the minimum MTU on the link, ensure that you configure the same MTU on the interfaces on both sides of the link.
- OSPFv3 drops the LS update if the packet size exceeds 8190. As a result, the neighbourship is terminated. Hence, configure the switch with the "ospfv3 MTU-ignore" command to avoid neighbourship termination.

OSPFv2 configuration

This section describes the tasks involved in configuring an OSPFv2 routing process. For a device using virtual routing, you can configure OSPFv2 for global as well as for user-defined virtual routers.

Configure OSPF areas, ranges, and virtual links

Complete the configuration of OSPF by defining areas, assigning route summarization ranges, and creating virtual links to connect discontiguous OSPF areas.

You can configure several OSPF area parameters, which include setting authentication, defining stub areas, and assigning specific costs to the default summary route. You can enable up to two OSPF process instances. Each OSPF process has its own associated areas and networks. Authentication provides password-based protection against unauthorized access to an area.

Stub areas are areas into which information on external routes is not sent. Instead, there is a default external route generated by the ABR into the stub area for destinations outside the autonomous system. To take advantage of the OSPF stub area support, default routing must be used in the stub area.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

For a virtual-router-aware device, choose the virtual router from the virtual routers drop-down list.

Step 3 Click **OSPF** and configure the following on the **OSPF** page:

- Check the check box of **Process 1**. You can enable up to two OSPF process instances for each context /virtual router . You must choose an OSPF process to be able to configure the Area parameters.

  If the device is using virtual routing, the ID fields display the unique process IDs generated for the chosen virtual router.

- Choose the **OSPF Role** from the drop-down list, and enter a description for it in the next field. The options are Internal, ABR, ASBR, and ABR & ASBR. See OSPF routing, on page 1447 for a description of the OSPF roles.

Step 4 Select **Area** > **Add** .

You can click **Edit** ( ) , or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 5 Configure the following area options for each OSPF process:

- **OSPF Process** — Choose the process ID. For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Area ID** —Designation of the area for which routes are to be summarized.
- **Area Type** —Choose one of the following:
  - **Normal** —(Default) Standard OSPF area.
  - **Stub** —A stub area does not have any routers or areas beyond it. Stub areas prevent Autonomous System (AS) External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you can prevent summary LSAs (Types 3 and 4) from being flooded into the area by not checking the **Summary Stub** check box.
  - **NSSA** —Makes the area a not-so-stubby area (NSSA). NSSAs accept Type 7 LSAs. You can disable route redistribution by not checking the **Redistribute** check box and checking the **Default Information Originate** check box. You can prevent summary LSAs from being flooded into the area by not checking the **Summary NSSA** check box.
- **Metric Value** —The metric used for generating the default route. The default value is 10. Valid metric values range from 0 through 16777214.
- **Metric Type** —The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Available Network** —Choose one of the available networks and click **Add** , or click **Add** ( ) to add a new network object. See Networks, on page 2982 for the procedure for adding networks.

- **Authentication** —Choose the OSPF authentication:
  - **None** —(Default) Disables OSPF area authentication.
  - **Password** —Provides a clear text password for area authentication, which is not recommended where security is a concern.
  - **MD5** —Allows MD5 authentication.
- **Default Cost** —The default cost for the OSPF area, which is used to determine the shortest paths to the destination. Valid values range from 0 through 65535. The default value is 1.
- Click **OK** to save the area configuration.

Step 6 Select **Range** > **Add** .

- Choose one of the available networks and whether to advertise, or,
- Click **Add** (+) to add a new network object. See Networks, on page 2982 for the procedure for adding networks.
- Click **OK** to save the range configuration.

Step 7 Select **Virtual Link** , click **Add** (+) , and configure the following options for each OSPF process:

- **Peer Router** —Choose the IP address of the peer router. To add a new peer router, click **Add** (+) . See Networks, on page 2982 for the procedure for adding networks.
- **Hello Interval** —The time in seconds between the hello packets sent on an interface. The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all routers and access servers on a specific network. Valid values range from 1 through 65535. The default is 10.

  The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.
- **Transmit Delay** —The estimated time in seconds that is required to send an LSA packet on the interface. The integer value must be greater than zero. Valid values range from 1 through 8192. The default is 1.

  LSAs in the update packet have their own ages incremented by this amount before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.
- **Retransmit Interval** —The time in seconds between LSA retransmissions for adjacencies that belong to the interface. The retransmit interval is the expected round-trip delay between any two routers on the attached network. The value must be greater than the expected round-trip delay, and can range from 1 through 65535. The default is 5.

  When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.
- **Dead Interval** —The time in seconds that hello packets are not seen before a neighbor indicates that the router is down. The dead interval is an unsigned integer. The default is four times the hello interval, or

40 seconds. The value must be the same for all routers and access servers that are attached to a common network. Valid values range from 1 through 65535.

- **Authentication** —Choose the OSPF virtual link authentication from the following:
  - **None** —(Default) Disables virtual link area authentication.
  - **Area Authentication** —Enables area authentication using MD5. Click **Add** , and enter the key ID, key, confirm the key, and then click **OK** .
  - **Password** —Provides a clear text password for virtual link authentication, which is not recommended where security is a concern.
  - **MD5** —Allows MD5 authentication. Click **Add** , and enter the key ID, key, confirm the key, and then click **OK** .

    Note
    Ensure to enter only numbers as the MD5 key ID.
  - **Key Chain** —Allows key chain authentication. Click **Add** , and create the key chain, and then click **Save** . For detailed procedure, see Create a key chain object, on page 2981 . Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency.
  - Click **OK** to save the virtual link configuration

Step 8 Click **Save** on the Routing page to save your changes.

OSPF is configured with designated areas, summarized ranges, and any necessary virtual links. Routing updates are generated based on your configuration.

What to do next

Continue with Configure OSPF Redistribution.

Configure OSPF redistribution

This task enables you to control the redistribution of routes between the OSPF routing processes on your Firewall Threat Defense device.

The Firewall Threat Defense device can control the redistribution of routes between the OSPF routing processes. The rules for redistributing routes from one routing process into an OSPF routing process are displayed. You can redistribute routes discovered by EIGRP, RIP and BGP into the OSPF routing process. You can also redistribute static and connected routes into the OSPF routing process.

Before you begin

Follow these steps to configure OSPF redistribution:

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 From **OSPF Role** drop-down, choose role.

Step 6 Click **Redistribution** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 7 Configure the following redistribution options for each OSPF process:

- **OSPF Process**—Choose the process ID. For a device using virtual routing, this drop-down list displays the unique process IDs generated for the selected virtual router.
- **Route Type**—Choose one of the following types:
  - **Static**—Redistributes static routes to the OSPF routing process.
  - **Connected**—Redistributes connected routes (routes established automatically by virtue of having the IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed AS external to the device. You can select whether to use subnets under the Optional list.
  - **OSPF**—Redistributes routes from another OSPF routing process, for example, internal, external 1 and 2, NSSA external 1 and 2, or whether to use subnets. You can select these options under the Optional list.
  - **BGP**—Redistribute routes from the BGP routing process. Add the AS number and whether to use subnets.
  - **RIP**—Redistributes routes from the RIP routing process. You can select whether to use subnets under the Optional list.

    Note
    As a user-defined virtual router does not support RIP, you cannot redistribute routes from RIP.

  - **EIGRP**—Redistribute routes from the EIGRP routing process. Add the AS number and whether to use subnets.
- **Metric Value**—Metric value for the routes being distributed. The default value is 10. Valid values range from 0 to 16777214.

  When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Tag Value**—Tag specifies the 32-bit decimal value attached to each external route that is not used by OSPF itself, but which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.
- **RouteMap**—Checks for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified,

but no route map tags are listed, no routes are imported. Or you can add a new route map by clicking **Add** (+). See Configure Route Map Entry to add a new route map.

Step 8 Click **OK** to save the redistribution configuration.

Step 9 Click **Save** on the Routing page to save your changes.

OSPF redistribution is now configured for the selected OSPF process with the specified route types and parameters.

What to do next

Continue with Configure OSPF Inter-Area filtering, on page 1459.

Configure OSPF Inter-Area filtering

Configure ABR type 3 LSA filtering to control which prefixes are sent between different OSPF areas and restrict other prefixes to their local OSPF area.

ABR type 3 LSA filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.

Before you begin

Follow these steps to configure OSPF inter-area filtering:

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **InterArea** > **Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete inter-areas.

Step 6 Configure these inter-area filtering options for each OSPF process:

- **OSPF Process**—For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.

- **Area ID**—The area for which routes are to be summarized.
- **PrefixList**—The name of the prefix. To add a new prefix list object, see Step 5.
- **Traffic Direction**—Inbound or outbound. Choose Inbound to filter LSAs coming into an OSPF area, or Outbound to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.

Step 7 Click **Add** (+), and enter a name for the new prefix list, and whether to allow overrides.

You must configure a prefix list before you can configure a prefix rule.

Step 8 Click **Add** to configure prefix rules, and configure these parameters:

- **Action**—Select **Block** or **Allow** for the redistribution access.
- **Sequence No**—The routing sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.
- **IP Address**—Specify the prefix number in the format of IP address/mask length.
- **Min Prefix Length**—(Optional) The minimum prefix length.
- **Max Prefix Length**—(Optional) The maximum prefix length.
- Click **OK** to save the inter-area filtering configuration.

Step 9 Click **Save** on the Routing page to save your changes.

The OSPF inter-area filtering configuration is saved and applied to control LSA distribution between the specified OSPF areas according to the configured prefix lists and traffic direction settings.

What to do next

Continue with Configure OSPF filter rules, on page 1460.

Configure OSPF filter rules

Configure OSPF filter rules to improve control of route distribution between OSPF areas by allowing only specified prefixes to be sent from one area to another area while restricting all other prefixes.

You can configure ABR Type 3 LSA filters for each OSPF process. ABR Type 3 LSA filters allow only specified prefixes to be sent from one area to another area and restrict all other prefixes. You can apply this type of area filtering out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF area at the same time. OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Filter Rule** > **Add**.

You can click **Edit** ( ), or use the right-click menu to cut, copy, past, insert, and delete filter rules.

Step 6 Configure the following filter rule options for each OSPF process:

- **OSPF Process**— For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Access List**—The access list for this OSPF process. To add a new standard access list object, click **Add** ( ) and see Configure a standard ACL object, on page 2953.
- **Traffic Direction**—Choose In or Out for the traffic direction being filtered. Choose In to filter LSAs coming into an OSPF area, or Out to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.
- **Interface**—The interface for this filter rule.

Step 7 Click **OK** to save the filter rule configuration.

Step 8 Click **Save** on the Routing page to save your changes.

The OSPF filter rules are configured and saved, enabling controlled route distribution between OSPF areas based on the specified access list and traffic direction settings.

What to do next

Continue with Configure OSPF summary addresses, on page 1461.

Configure OSPF summary addresses

Configure OSPF summary addresses to decrease the size of the OSPF link-state database by advertising a single route for all redistributed routes that are included for a specified network address and mask.

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the Firewall Threat Defense device to advertise a single route for all the redistributed routes that are included for a specified network address and mask. This configuration decreases the size of the OSPF link-state database. Routes that match the specified IP address mask pair can be suppressed. The tag value can be used as a match value for controlling redistribution through route maps.

Routes learned from other routing protocols can be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. Summary routes help reduce the size of the routing table.

Using summary routes for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address. Only routes from other routing protocols that are being redistributed into OSPF can be summarized.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Summary Address** > **Add**.

You can click **Edit** ( ) to edit, or use the right-click menu to cut, copy, past, insert, and delete summary addresses.

Step 6 Configure the following summary address options for each OSPF process:

- **OSPF Process**— For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Available Network**—The IP address of the summary address. Select one from the Available networks list and click **Add**, or to add a new network, click **Add** ( ). See Networks, on page 2982 for the procedure for adding networks.
- **Tag**—A 32-bit decimal value that is attached to each external route. This value is not used by OSPF itself, but may be used to communicate information between ASBRs.
- **Advertise—Advertises** the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Step 7 Click **OK** to save the summary address configuration.

Step 8 Click **Save** on the Routing page to save your changes.

The OSPF summary address configuration is saved and the device will advertise a single route for all redistributed routes that are included for the specified network address and mask.

What to do next

Continue with Configure OSPF interfaces and neighbors, on page 1462.

Configure OSPF interfaces and neighbors

You can change some interface-specific OSPFv2 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval, the dead interval, and the authentication key. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

You need to define static OSPFv2 neighbors to advertise OSPFv2 routes over a point-to-point, non-broadcast network. This feature lets you broadcast OSPFv2 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Interface** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 6 Configure the following Interface options for each OSPF process:

- **Interface**—The interface you are configuring.

  Note
  If the device is using virtual routing, this drop-down list displays only those interfaces that belong to the router.

- **Default Cost**—The cost of sending a packet through the interface. The default value is 10.
- **Priority**—The designated router for a network. Valid values range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router.

  When two routers connect to a network, both attempt to become the designated router. The device with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router. This setting does not apply to interfaces that are configured as point-to-point interfaces.

- **MTU Ignore—OSPF** checks whether neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the IP MTU configured on the incoming interface, OSPF adjacency is not established.
- **Database Filter**—Use this setting to filter the outgoing LSA interface during synchronization and flooding. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives. In a fully meshed topology, this flooding can waste bandwidth and lead to excessive link and CPU usage. Checking this check box prevents OSPF flooding of the LSA on the selected interface.
- **Hello Interval**—Specifies the interval, in seconds, between hello packets sent on an interface. Valid values range 1–8192 seconds. The default value is 10 seconds.

  The smaller the hello interval, the faster topological changes are detected, but more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface.

- **Transmit Delay**—Estimated time in seconds to send an LSA packet on the interface. Valid values range 1–65535 seconds. The default is 1 second.

  LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

- **Retransmit Interval**—Time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.

  When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.

- **Dead Interval**—Time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range 1–65535.
- **Hello Multiplier**—Specifies the number of Hello packets to be sent per second. Valid values are 3–20.
- **Point-to-Point**—Lets you transmit OSPF routes over VPN tunnels.
- **Authentication**—Choose the OSPF interface authentication from the following:
  - **None**—(Default) Disables interface authentication.
  - **Area Authentication**—Enables interface authentication using MD5. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.
  - **Password**—Provides a clear text password for virtual link authentication, which is not recommended where security is a concern.
  - **MD5**—Allows MD5 authentication. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.

    Note
    Enter only letters, numbers, and special characters for the MD5 key ID. Do not enter spaces; spaces are truncated in management center.

  - **Key Chain**—Allows key chain authentication. Click **Add**, and create the key chain, and then click **Save**. For detailed procedure, see Create a key chain object, on page 2981. Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency.
- **Enable BFD**—Allows you to enable BFD on this interface.
- **Enter Password**—The password you configure if you choose Password as the type of authentication.
- **Confirm Password**—Confirm the password that you chose.

Step 7 Select **Neighbor** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 8 Configure the following parameters for each OSPF process:

- **OSPF Process**—Choose 1 or 2.
- **Neighbor**—Choose one of the neighbors in the drop-down list, or click **Add** () to add a new neighbor; enter the name, description, network, whether to allow overrides, and then click **Save**.
- **Interface**—Choose the interface associated with the neighbor.

Step 9 Click **OK** and click **Save** to save the changes.

OSPF interface and neighbor configuration is completed. The configured parameters are applied to the selected interfaces and neighbors, enabling OSPF routing functionality on the specified network segments.

Configure OSPF Advanced Properties

Enable advanced OSPF features for optimized routing behavior and enhanced protocol reliability.

The Advanced Properties allows you to configure options, such as syslog message generation, administrative route distances, an LSA timer, and graceful restarts.

Graceful Restarts

The Firewall Threat Defense device may experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a scheduled hitless software upgrade. You can configure graceful restart on OSPFv2 by using either using NSF Cisco (RFC 4811 and RFC 4812) or NSF IETF (RFC 3623).

Note NSF capability is also useful in HA mode and clustering.

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing** .

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF** > **Advanced** .

Step 5 Select **General**, and configure the following: .

- **Router ID**: Choose Automatic or IP Address (appears for non-cluster and a cluster in spanned etherchannel mode) or Cluster Pool (appears for a cluster in individual interface mode) for the router ID.

If you choose IP address, enter the IP address in the adjacent field. If you choose Cluster Pool, choose the IPv4 cluster pool value in the adjacent drop-down field. For information on creating the cluster pool address, see Configure address pools, on page 2954 .

- **Ignore LSA MOSPF**: Suppresses syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets.
- **RFC 1583 Compatible**: Configures RFC 1583 compatibility as the method used to calculate summary route costs. Routing loops can occur with RFC 1583 compatibility enabled. Disable it to prevent routing loops. All OSPF routers in an OSPF routing domain should have RFC compatibility set identically.
- **Adjacency Changes**: Defines the adjacency changes that cause syslog messages to be sent.

  By default, a syslog message is generated when an OSPF neighbor goes up or down. You can configure the router to send a syslog message when an OSPF neighbor goes down and also a syslog for each state.

  - **Log Adjacency Changes**: Causes the Firewall Threat Defense device to send a syslog message whenever an OSPF neighbor goes up or down. This setting is checked by default.
  - **Log Adjacency Change Details**: Causes the Firewall Threat Defense device to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.

- **Administrative Route Distance**: Allows you to modify the settings that were used to configure administrative route distances for **inter-area**, **intra-area**, and **external** IPv6 routes. The administrative route distance is an integer from 1 to 254. The default is 110.
- **LSA Group Pacing**: Specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.
- **Enable Default Information Originate**: Check the **Enable** check box to generate a default external route into an OSPF routing domain and configure the following options:
  - **Always advertise the default route**: Ensures that the default route is always advertised.
  - **Metric Value**: Metric used for generating the default route. Valid metric values range from 0 to 16777214. The default value is 10.
  - **Metric Type**: The external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are 1 (Type 1 external route) and 2 (Type 2 external route). The default is Type 2 external route.
  - **RouteMap**: Choose the routing process that generates the default route if the route map is satisfied or click **Add** (+) to add a new one. See Configure Route Map Entry to add a new route map.

Step 6 Click **OK** to save the general configuration.

Step 7 Select **Non Stop Forwarding**, and configure Cisco NSF graceful restart for OSPFv2, for an NSF-capable or NSF-aware device.

Note
There are two graceful restart mechanisms for OSPFv2, Cisco NSF and IETF NSF. Only one of these graceful restart mechanisms can be configured at a time for an OSPF instance. An NSF-aware device can be configured as both Cisco NSF helper and IETF NSF helper but a NSF-capable device can be configured in either Cisco NSF or IETF NSF mode at a time for an OSPF instance.

a) Check the **Enable Cisco Non Stop Forwarding Capability** check box.
b) (Optional) Check the **Cancel NSF restart when non-NSF-aware neighboring networking devices are detected** check box if required.
c) (Optional) Make sure the **Enable Cisco Non Stop Forwarding Helper** mode check box is unchecked to disable the helper mode on an NSF-aware device.

Step 8 Configure IETF NSF Graceful Restart for OSPFv2, for an NSF-capable or NSF-aware device: .

a) Check the **Enable IETF Non Stop Forwarding Capability** check box.
b) In the **Length of graceful restart interval (seconds)** field, enter the restart interval in seconds. The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.
c) (Optional) Make sure the **Enable IETF nonstop forwarding (NSF) for helper mode** check box is unchecked to disable the IETF NSF helper mode on an NSF-aware device.
d) **Enable Strict Link State advertisement checking**: When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.
e) **Enable IETF Non Stop Forwarding**: Enables non stop forwarding, which allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. OSPF uses extensions to the OSPF protocol to recover its state from neighboring OSPF devices. For the recovery to work, the neighbors must support the NSF protocol extensions and be willing to act as "helpers" to the device that is restarting. The neighbors must also continue forwarding data traffic to the device that is restarting while protocol state recovery takes place.

Advanced OSPF settings are applied. The device now uses customized OSPF general and graceful restart properties for improved routing behavior, syslog management, and protocol resilience.

OSPFv3 configuration

OSPFv3 configuration is a routing process setup that

- enables IPv6 routing functionality on network devices,
- involves multiple configuration tasks to establish the routing process, and
- supports configuration only for global virtual routers in virtual routing environments.

Virtual routing limitations

For a device using virtual routing, you can configure OSPFv3 only for its global virtual router and not for its user-defined virtual router.

Configure OSPFv3 Areas, Route Summaries, and Virtual Links

-

To enable OSPFv3, you need to create an OSPFv3 routing process, create an area for OSPFv3, enable an interface for OSPFv3, and then redistribute the route into the targeted OSPFv3 routing process.

Before you begin

-

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click the **Routing** tab. On the tab's left pane, click **OSPFv3** .

Step 3 By default **Enable Process 1** is selected. You can enable up to two OSPF process instances.

Step 4 Choose the OSPFv3 role from the drop-down list, and enter a description for it. The options are Internal, ABR, ASBR, and ABR and ASBR. See OSPF routing, on page 1447 for descriptions of the OSPFv3 roles.

Step 5 Select **Area** > **Add** .

You can click **Edit** ( ) , or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 6 Select **General** , and configure the following options for each OSPF process:

- **Area ID** —The area for which routes are to be summarized.
- **Cost** —The metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.
- **Type** —Specifies Normal, NSSA, or Stub. If you select Normal, there are no other parameters to configure. If you select Stub, you can choose to send summary LSAs in the area. If you select NSSA, you can configure the next three options:
  - **Allow Sending summary LSA into this area** —Allows the sending of summary LSAs into the area.
  - **Imports routes to normal and NSSA area** —Allows redistribution to import routes to normal and not to stubby areas.
  - **Defaults information originate** —Generates a default external route into an OSPFv3 routing domain.
- **Metric** —Metric used for generating the default route. The default value is 10. Valid metric values range from 0 to 16777214.
- **Metric Type** —The metric type is the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.

Step 7 Click **OK** to save the general configuration.

Step 8 (Not applicable for Internal OSPFv3 Role) Select **Route Summary** > **Add Route Summary** .

You can click **Edit** ( ) , or use the right-click menu to cut, copy, past, insert, and delete route summaries.

Step 9 Configure the following route summary options for each OSPF process:

- **IPv6 Prefix/Length** —The IPv6 prefix. To add a new network object, click **Add** ( ) . See Networks, on page 2982 for the procedure for adding networks.
- **Cost** —The metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.

- **Advertise** —Advertises the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Step 10 Click **OK** to save the route summary configuration.

Step 11 (Not applicable for Internal OSPFv3 Role) Select **Virtual Link** , click **Add Virtual Link** , and configure the following options for each OSPF process:

- **Peer RouterID** —Choose the IP address of the peer router. To add a new network object, click **Add** (+) . See Networks, on page 2982 for the procedure for adding networks.
- **TTL Security** —Enables TTL security check. The value for the hop-count is a number from 1 to 254. The default is 1.

  OSPF sends outgoing packets with an IP header Time to Live (TTL) value of 255 and discards incoming packets that have TTL values less than a configurable threshold. Because each device that forwards an IP packet decrements the TTL, packets received via a direct (one-hop) connection have a value of 255. Packets that cross two hops have a value of 254, and so on. The receive threshold is configured in terms of the maximum number of hops that a packet may have traveled.
- **Dead Interval** —The time in seconds that hello packets are not seen before a neighbor indicates that the router is down. The default is four times the hello interval, or 40 seconds. Valid values range from 1 to 65535.

  The dead interval is an unsigned integer. The value must be the same for all routers and access servers that are attached to a common network.
- **Hello Interval** —The time in seconds between the hello packets sent on an interface. Valid values range from 1 to 65535. The default is 10.

  The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all routers and access servers on a specific network. The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.
- **Retransmit Interval** —The time in seconds between LSA retransmissions for adjacencies that belong to the interface. The retransmit interval is the expected round-trip delay between any two routers on the attached network. The value must be greater than the expected round-trip delay, and can range from 1 to 65535. The default is 5.

  When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.
- **Transmit Delay** —The estimated time in seconds that is required to send an LSA packet on the interface. The integer value must be greater than zero. Valid values range from 1 to 8192. The default is 1.

  LSAs in the update packet have their own ages incremented by this amount before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

Step 12 Click **OK** to save the virtual link configuration.

Step 13 Click **Save** on the Router page to save your changes.

What to do next

Continue with Configure OSPFv3 Redistribution .

Configure OSPFv3 redistribution

Configure OSPFv3 redistribution to control how routes are redistributed between OSPF routing processes and other routing protocols on the Secure Firewall Threat Defense device.

The Secure Firewall Threat Defense device can control the redistribution of routes between the OSPF routing processes. The rules for redistributing routes from one routing process into an OSPF routing process are displayed. You can redistribute routes discovered by EIGRP, RIP and BGP into the OSPF routing process. You can also redistribute static and connected routes into the OSPF routing process.

Before you begin

Follow these steps to configure OSPFv3 redistribution:

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** > **OSPF**.

Step 3 Select **Redistribution**, and click **Add**.

You can click **Edit** ( ), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 4 Configure the following redistribution options for each OSPF process:

- **Source Protocol**—The source protocol from which routes are being redistributed. The supported protocols are connected, OSPF, Static, EIGRP, and BGP. If you choose OSPF, you must enter the Process ID in the **Process ID** field. If you choose BGP, you must add the AS number in the **AS Number** field.
- **Metric**—Metric value for the routes being distributed. The default value is 10. Valid values range from 0 to 16777214.

  When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Tag**—Tag specifies the 32-bit decimal value attached to each external route that is not used by OSPF itself, but which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.
- **Route Map**—Checks for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified, but no route map tags are listed, no routes are imported. Or you can add a new route map by clicking **Add** ( ). See Configure a route map, on page 3017 for the procedure to add a new route map.

- **Process ID**—The OSPF process ID, either 1 or 2.

  Note
  The Process ID is enabled the OSPFv3 process is redistributing a route learned by another OSPFv3 process.

- **Match**—Enables OSPF routes to be redistributed into other routing domains:
  - **Internal** for routes that are internal to a specific autonomous system.
  - **External 1**for routes that are external to the autonomous system, but are imported into OSPFv3 AS Type 1 external routes.
  - **External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 AS Type 2 external routes.
  - **NSSA External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 AS Type 1 external routes.
  - **NSSA External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 AS Type 2 external routes.

Step 5 Click **OK** to save the redistribution configuration.

Step 6 Click **Save** on the Routing page to save your changes.

The OSPFv3 redistribution configuration is saved and routes are redistributed according to the specified parameters.

What to do next

Continue with Configure OSPFv3 summary prefixes, on page 1471.

Configure OSPFv3 summary prefixes

Configure OSPFv3 summary prefixes to control route advertisement and suppression based on IPv6 prefix and mask pairs.

You can configure the Firewall Threat Defense device to advertise routes that match a specified IPv6 prefix and mask pair.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** > **OSPFv3**.

Step 3 Select **Summary Prefix** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete summary prefixes.

Step 4 Configure the following summary prefix options for each OSPF process:

- **IPv6 Prefix/Length**—The IPv6 prefix and prefix length label. Select one from the list or click **Add** (+) to add a new network object. See Networks, on page 2982 for the procedure for adding networks.
- **Advertise**— Advertises routes that match the specified prefix and mask pair. Uncheck this check box to suppress routes that match the specified prefix and mask pair.
- (Optional) **Tag**—A value that you can use as a match value for controlling redistribution through route maps.

Step 5 Click **OK** to save the summary prefix configuration.

Step 6 Click **Save** on the Routing page to save your changes.

The OSPFv3 summary prefix configuration is saved and will control route advertisement based on your specified settings.

What to do next

Continue with Configure OSPFv3 interfaces, authentication, and neighbors, on page 1472.

Configure OSPFv3 interfaces, authentication, and neighbors

Set up OSPFv3 routing, authentication, and neighbor relationships to support IPv6 dynamic network routing.

You can change certain interface-specific OSPFv3 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval and the dead interval. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

For a successful implementation of OSPFv3 authentication with Nexus switches, ensure that you have the compatible version of the switch, for example, Nexus 3000, 7000, and 9000 series switches.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** > **OSPFv3**.

Step 3 Select **Interface** > **Add**.

You can click **Edit** to edit, or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 4 Configure the following interface options for each OSPFv3 process:

- **Interface** —The interface you are configuring.
- **Enable OSPFv3** —Enables OSPFv3.
- **OSPF Process** —Choose 1 or 2.
- **Area** —The area ID for this process.

- **Instance** —Specifies the area instance ID to be assigned to the interface. An interface can have only one OSPFv3 area. You can use the same area on multiple interfaces, and each interface can use a different area instance ID.

Step 5 Select **Properties** , and configuring the following options for each OSPFv3 process:

- **Filter Outgoing Link Status Advertisements** —Filters outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.
- **Disable MTU mismatch detection** —Disables the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.
- **Flood Reduction** —Changes normal LSAs into Do Not Age LSAs, so that they don't get flooded every 3600 seconds across areas.

  OSPF LSAs are refreshed every 3600 seconds. In large OSPF networks, this can lead to large amounts of unnecessary LSA flooding from area to area.
- **Point-to-Point Network** —Lets you transmit OSPF routes over VPN tunnels. When an interface is configured as point-to-point, non-broadcast, the following restrictions apply:
  - You can define only one neighbor for the interface.
  - You need to manually configure the neighbor.
  - You need to define a static route pointing to the crypto endpoint.
  - If OSPF over a tunnel is running on the interface, regular OSPF with an upstream router cannot be run on the same interface.
  - You should bind the crypto map to the interface before specifying the OSPF neighbor to ensure that the OSPF updates are passed through the VPN tunnel. If you bind the crypto map to the interface after specifying the OSPF neighbor, use the **clear local-host all** command to clear OSPF connections so that the OSPF adjacencies can be established over the VPN tunnel.
- **Broadcast** — Specifies that the interface is a broadcast interface. By default, this check box is checked for Ethernet interfaces. Uncheck this check box to designate the interface as a point-to-point, nonbroadcast interface. Specifying an interface as point-to-point, nonbroadcast lets you transmit OSPF routes over VPN tunnels.
- **Cost** —Specifies the cost of sending a packet on the interface. Valid values for this setting range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router. This setting does not apply to interfaces that are configured as point-to-point, nonbroadcast interfaces.

  When two routers connect to a network, both attempt to become the designated router. The device with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router.
- **Priority** —Determines the designated router for a network. Valid values range from 0 to 255.
- **Dead Interval** —Time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range from 1 to 65535.
- **Hello Interval** — Time period in seconds between OSPF packets that the router will send before adjacency is established with a neighbor. Once the routing device detects an active neighbor, the hello packet interval

changes from the time specified in the poll interval to the time specified in the hello interval. Valid values range from 1 to 65535 seconds.

- **Retransmit Interval** —Time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.
- **Transmit Delay** —Estimated time in seconds to send a link-state update packet on the interface. Valid values range from 1 to 65535 seconds. The default is 1 second.
- **Enable BFD** —Allows you to enable BFD on this interface.

Step 6 Select **Authentication** , and configure the following options for each OSPFv3 process:

- **Type** —Type of authentication. The available options are Area, Interface, and None. The None option indicates that no authentication is used.
- **Security Parameters Index** — A number from 256 to 4294967295. Configure this if you chose Interface as the type.
- **Authentication** —Type of authentication algorithm. Supported values are SHA-1 and MD5. Configure this if you chose Interface as the type.
- **Authentication Key** — When MD5 authentication is used, the key must be 32 hexadecimal digits (16 bytes) long. When SHA-1 authentication is used, the key must be 40 hexadecimal digits (20 bytes) long.
- **Encrypt Authentication Key** —Enables encryption of the authentication key.
- **Include Encryption** — Enables encryption.
- **Encryption Algorithm** —Type of encryption algorithm. Supported value is DES. The NULL entry indicates no encryption. Configure this if you chose **Include Encryption** .
- **Encryption Key** —Enter the encryption key. Configure this if you chose **Include Encryption** .
- **Encrypt Key** —Enables the key to be encrypted.

Step 7 Click **OK** to save the authentication configuration.

Step 8 Select **Neighbor** , click **Add** , and configure the following options for each OSPFv3 process:

- **Link Local Address** —The IPv6 address of the static neighbor.
- **Cost** —Enables cost. Enter the cost in the **Cost** field, and check the **Filter Outgoing Link State Advertisements** if you want to advertise.
- (Optional) **Poll Interval** —Enables the poll interval. Enter the **Priority** level and the **Poll Interval** in seconds.
- Click **Add** to add the neighbor.

Step 9 Click **OK** to save the Interface configuration.

OSPFv3 interface, authentication, and neighbor settings are configured on the device. OSPFv3 processes will operate according to the specified parameters.

Configure OSPFv3 Advanced Properties

Enable and fine-tune OSPFv3 advanced properties to optimize routing, ensure stability, and automate operations for your device

The Advanced Properties allows you to configure options, such as syslog message generation, administrative route distances, passive OSPFv3 routing, LSA timers, and graceful restarts.

Graceful Restarts

The Firewall Threat Defense device may experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a scheduled hitless software upgrade. You can configure graceful restart on OSPFv3 using graceful-restart (RFC 5187).



Note NSF capability is also useful in HA mode and clustering.

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **OSPFv3** > **Advanced** .

Step 3 For **Router ID** , choose Automatic or IP Address (appears for non-cluster and a cluster in spanned etherchannel mode) or Cluster Pool (appears for a cluster in individual interface mode) . If you choose IP Address, enter the IPv6 address in the **IP Address** field. If you choose Cluster Pool, choose the IPv6 cluster pool value from the **Cluster Pool** down-down field. For information on creating the cluster pool address, see Configure address pools, on page 2954 .

Step 4 Check the **Ignore LSA MOSPF** check box if you want to suppress syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets.

Step 5 Select **General** , and configure the following:

- **Adjacency Changes** —Defines the adjacency changes that cause syslog messages to be sent.

  By default, a syslog message is generated when an OSPF neighbor goes up or down. You can configure the router to send a syslog message when an OSPF neighbor goes down and also a syslog for each state.

- **Adjacency Changes** —Causes the Firewall Threat Defense device to send a syslog message whenever an OSPF neighbor goes up or down. This setting is checked by default.
- **Include Details** —Causes the Firewall Threat Defense device to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.

- **Administrative Route Distances** —Allows you to modify the settings that were used to configure administrative route distances for inter-area, intra-area, and external IPv6 routes. The administrative route distance is an integer from 1 to 254. The default is 110.
- **Default Information Originate** —Check the **Enable** check box to generate a default external route into an OSPFv3 routing domain and configure the following options:
  - **Always Advertise** —Will always advertise the default route whether or not one exists.
  - **Metric** —Metric used for generating the default route. Valid metric values range from 0 to 16777214. The default value is 10.
  - **Metric Type** —The external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are 1 (Type 1 external route) and 2 (Type 2 external route). The default is Type 2 external route.
  - **Route Map** —Choose the routing process that generates the default route if the route map is satisfied or click **Add** (+) to add a new one. See Configure a route map, on page 3017 to add a new route map.

Step 6 Click **OK** to save the general configuration.

Step 7 Select **Passive Interface**, select the interfaces on which you want to enable passive OSPFv3 routing from the Available Interfaces list, and click **Add** to move them to the Selected Interfaces list.

Passive routing assists in controlling the advertisement of OSPFv3 routing information and disables the sending and receiving of OSPFv3 routing updates on an interface.

Step 8 Click **OK** to save the passive interface configuration.

Step 9 Select **Timer**, and configure the following LSA pacing and SPF calculation timers:

- **Arrival**—Specifies the minimum delay in milliseconds that must pass between acceptance of the same LSA arriving from neighbors. The range is from 0 to 6000,000 milliseconds. The default is 1000 milliseconds.
- **Flood Pacing**—Specifies the time in milliseconds at which LSAs in the flooding queue are paced in between updates. The configurable range is from 5 to 100 milliseconds. The default value is 33 milliseconds.
- **Group Pacing**—Specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.
- **Retransmission Pacing**—Specifies the time in milliseconds at which LSAs in the retransmission queue are paced. The configurable range is from 5 to 200 milliseconds. The default value is 66 milliseconds.
- **LSA Throttle** —Specifics the delay in milliseconds to generate the first occurrence of the LSA. The default value is 0 millisecond. The minimum specifies the minimum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds. The maximum specifies the maximum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds.

Note
For LSA throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

- **SPF Throttle** —Specifies the delay in milliseconds to receive a change to the SPF calculation. The default value is 5000 milliseconds. The minimum specifies the delay in milliseconds between the first and second SPF calculations. The default value is 10000 milliseconds. The maximum specifies the maximum wait time in milliseconds for SPF calculations. The default value is 10000 milliseconds.

  Note
  For SPF throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

Step 10 Click **OK** to save the LSA timer configuration.

Step 11 Select **Non Stop Forwarding** , and check the **Enable graceful-restart helper** check box. This is checked by default. Uncheck this to disable the graceful-restart helper mode on an NSF-aware device.

Step 12 Check the **Enable link state advertisement** check box to enable strict link state advertisement checking.

When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.

Step 13 Check the **Enable graceful-restart (Use when Spanned Cluster or Failover Configured)** and enter the graceful-restart interval in seconds. The range is 1-1800. The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.

Step 14 Click **OK** to save the graceful restart configuration.

Step 15 Click **Save** on the Routing page to save your changes.

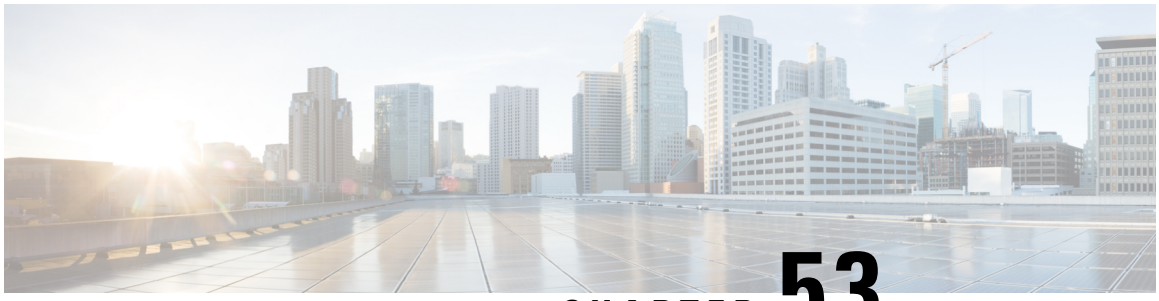
The device is configured with your OSPFv3 advanced properties. All changes are applied, enhancing routing resilience, logging, and failover handling.

History for OSPF

This reference provides a summary of the key features introduced in OSPF and the versions in which they became available, along with details about their implementation.

Table 94: Feature history for OSPF

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
BFD Support for OSPF v2 and v3	7.4	7.4	You can enable BFD on OSPFv2 and OSPFv3 interfaces. New/modified screens: • **Configuration** > **Device Setup** > **Routing** > **OSPFv2** • **Configuration** > **Device Setup** > **Routing** > **OSPFv3**



CHAPTER 53

EIGRP

EIGRP is a routing protocol that

- enables routers to efficiently exchange information,
- maintain network topology tables, and
- ensures secure data path selection and communication across networks.

EIGRP Operational Overview

This section covers the configuration of the Firewall Threat Defense to perform these routing tasks:

- Route data traffic
- Perform authentication
- Redistribute routing information

- EIGRP routing, on page 1479
- Requirements and prerequisites for EIGRP, on page 1480
- Guidelines and limitations for EIGRP routing, on page 1481
- Configure EIGRP routing on firewall devices, on page 1482

EIGRP routing

An Enhanced Interior Gateway Routing Protocol (EIGRP) is a routing protocol that

- sends updates only when network topology changes, not at regular intervals,
- offers rapid convergence and supports variable-length subnet masks, and
- supports partial updates and multiple network layer protocols.

EIGRP capabilities and operations

EIGRP stores all neighbor routing tables for quick adaptation to alternate routes. If a suitable route does not exist, EIGRP queries its neighbors to discover an alternate route. It propagates these queries through the network to locate a route. EIGRP's support for variable-length subnet masks allows routes to be automatically

summarized on a network boundary. Additionally, EIGRP can be configured to summarize any bit boundary at any interface.

EIGRP sends partial updates when route metrics change, limiting propagation to necessary routers, which minimizes bandwidth use compared to protocols that send periodic full updates.

To learn about routers on attached networks, EIGRP uses neighbor discovery through multicast hello packets to announce their presence on the network. When a new neighbor is detected through hello packets, a topology table exchange occurs.

EIGRP hello packets are normally sent as multicast messages and do not require explicit replies. For manually configured static neighbors, hello packets are sent as unicast messages to the specified neighbor.

Once this neighbor relationship is established, routing updates are exchanged only on topology changes. The relationship persists through regular multicast hello packets, with devices expected to respond within advertised hold times. Hold time is the time within which threat defense can expect to receive a hello packet from that neighbor. If the device does not receive a hello packet within the advertised hold time, it considers the neighbor unavailable.

EIGRP employs mechanisms like neighbor discovery/recovery, Reliable Transport Protocol (RTP), and Diffusing Update Algorithm (DUAL) for route computations. DUAL retains all routes to a destination, selecting the least-cost route for packet forwarding while retaining others in case of network changes. If the main route fails, another route is chosen from the feasible successors. A successor is a neighboring router that is used for packet forwarding that has a least-cost path to a destination. DUAL utilizes a feasibility condition to ensure that the path is not part of a routing loop.

If a feasible successor is not found in the topology table, a route recomputation takes place. During route recomputation, DUAL queries the EIGRP neighbors for a route. The query is propagated to successive neighbors. If a feasible successor is not found, an unreachable message is returned.

During route recomputation, DUAL marks the route as active. By default, threat defense waits for three minutes to receive a response from its neighbors. If the device does not receive a response from a neighbor, the route is marked as stuck-in-active. All routes in the topology table that point to the unresponsive neighbor as a feasible successor are removed.

Requirements and prerequisites for EIGRP

This topic provides the necessary prerequisites and requirements for deploying EIGRP within a Cisco network environment.

Model support

Firewall Threat Defense

Firewall Threat Defense Virtual

Supported domains

Any

User roles

Admin

Network Admin

Guidelines and limitations for EIGRP routing

Firewall mode guidelines

Use routed firewall mode only for EIGRP routing.

Device guidelines

Configure only one EIGRP process per device.

- EIGRP can be configured through management center UI on Firewall Threat Defense version 6.6 and later.

Interface guidelines

Use only routed interfaces with logical names and IP addresses for EIGRP routing processes.

- EIGRP can only incorporate interfaces from the global virtual router. EIGRP can learn, filter, and redistribute routes across routing protocols in the global virtual router.
- Supports physical, EtherChannel, redundant, and subinterfaces only. However, the members of EtherChannel interfaces are not supported.
-
- Passive interfaces cannot be configured as neighbor interfaces.

IP address and network objects support

Use only IPv4 addresses and standard access list objects for EIGRP configuration.

- Ranges, FQDNs, and wildcard masks are not supported.

Redistribution guidelines

Configure route tagging to prevent routing loops when redistributing between EIGRP and OSPF.

- BGP, OSPF, and RIP in the global virtual router can redistribute routes toward EIGRP.
- EIGRP can redistribute to BGP, OSPF, RIP, Static, and Connected in the global virtual router.
- When EIGRP is configured on a device that is a part of an OSPF network or vice versa, ensure that the OSPF router is configured to tag the route (EIGRP does not support route tag).

  When redistributing EIGRP into OSPF and OSPF into EIGRP, a routing loop can occur when there is an outage on one of the links, interfaces, or even when the route originator is down. To prevent the redistribution of routes from one domain back into the same domain, a router can tag a route that belongs to a domain while it is redistributing, and those routes can be filtered on the remote router based on the same tag. Because the routes will not be installed into the routing table, they will not be redistributed back into the same domain.

Deployment process guidelines

Disable and redeploy EIGRP when changing the AS number to prevent deployment failures.

To effectively change the AS number of an existing EIGRP setup, you should disable and redeploy EIGRP configurations to avoid repeated deployments and ensure error-free deployment processes.

Upgrade guidelines

Migrate FlexConfig EIGRP policies to UI management after upgrading to version 7.2 and later.

When you upgrade to version 7.2 and later, and if the previous version includes any FlexConfig EIGRP policies, the management center displays a warning message during deployment. It does not stop the deployment process; a warning message is displayed. After deployment, to manage the EIGRP policies from the UI (**Devices** > **Device Management**, click the **Edit** (✎) icon. Then, click **Routing>EIGRP**), you must redo the configuration in the **EIGRP** page and remove the configuration from FlexConfig. To automate policy creation in the UI, Cloud-Delivered Firewall Management Center provides an option to migrate the policies from FlexConfig to the UI. For more information, refer to Migrate FlexConfig policies, on page 1352.

Configure EIGRP routing on firewall devices

Enable EIGRP on a firewall device for efficient routing within an autonomous system.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click the **Routing** tab.

Step 3 Under Global, click **EIGRP**.

Step 4 Check the **Enable EIGRP** check box to enable the EIGRP routing process.

Step 5 In the **AS Number** field, enter the autonomous system (AS) number for the EIGRP process.

Valid values are from 1 to 65535. The AS number identifies the local EIGRP routing domain and must match on EIGRP neighbors.

Step 6 To configure other EIGRP properties, refer to these topics for detailed steps:

a. Configure EIGRP settings, on page 1483.

b. Configure EIGRP neighbors settings, on page 1483.

c. Configure EIGRP filter rules, on page 1484.

d. Configure EIGRP redistribution settings, on page 1485.

e. Configure EIGRP summary address settings, on page 1486.

f. Configure EIGRP interfaces settings, on page 1487.

g. Configure EIGRP advanced settings, on page 1488.

The firewall device is now configured with the base EIGRP process and the designated AS number. It can participate in EIGRP routing after the required interfaces, networks, and related settings are configured.

Configure EIGRP settings

Use this procedure when you need to set up EIGRP routing protocol parameters on your device. EIGRP configuration includes enabling automatic network summarization, selecting networks that participate in the routing process, and configuring passive interfaces that do not exchange routing updates.

Procedure

Step 1 On the **EIGRP** page, click the **Setup** tab.

Step 2 Check the **Auto Summary** check box to enable EIGRP to summarize at network boundaries.

Note
Enabling Auto Summary can cause routing problems if you have noncontiguous networks.

Step 3 In the **Available Networks/Hosts** box, select the networks or hosts that should participate in the EIGRP routing process, and then click **Add**. To add a new network object, click **Add** (+). For detailed instructions on adding networks, refer to Networks, on page 2982.

Step 4 To configure passive interfaces, check the **Passive Interface** check box. In EIGRP, a passive interface does not send or receive routing updates.

a) To specify selective interfaces as passive, click the **Selected Interface** radio button. In the **Available Interfaces** box, select the interfaces, and click **Add**.

b) To specify all interfaces as passive, click the **All Interfaces** radio button.

Step 5 Click **Save**.

The EIGRP settings are configured and saved. The routing protocol will operate according to the specified configuration, including automatic summarization behavior, network participation, and passive interface settings.

Configure EIGRP neighbors settings

Define static neighbors for the EIGRP process. This results in unicast hello packets being sent to that neighbor.

Procedure

Step 1 On the **EIGRP** page, click the **Neighbors** tab.

Step 2 Click +**Add**.

Step 3 Select the interface through which the neighbor is available from the **Interface** drop-down list.

Step 4 Select the IP address of the static neighbor from the **Neighbor** drop-down list.

Step 5 To add a new network object, click +. For instructions on adding network objects, refer to Networks, on page 2982.

Step 6 Click **OK**.

Step 7 Click **Save**.

The EIGRP static neighbor is configured, and hello packets are sent as unicast to the specified neighbor through the designated interface.

Configure EIGRP filter rules

Filter rules provide granular control over EIGRP route advertisements and acceptances, allowing you to manage network traffic flow and routing table contents based on specific criteria such as interface or protocol type.

To configure EIGRP filter rules, perform these steps:

Procedure

Step 1 On the **EIGRP** page, click the **Filter Rules** tab.

Step 2 Click **Add**.

Step 3 In the **Add Filter Rules** dialog box, select the direction for the rule from the **Filter Direction** drop-down list:

- **Inbound**: The rule filters routes received in incoming EIGRP updates.
- **Outbound**: The rule filters routes advertised in outgoing EIGRP updates.

Step 4 Under **Type**, choose one of these options:

- To select the interface to which the filtering rule applies, click **Interface**, and then select the interface from the **Select Interface** drop-down list. Select the access list from the **Access List** drop-down list. The list defines the networks that are to be received and suppressed in routing updates. To add a new standard access list object, click +. For more information, refer to Configure a standard ACL object, on page 2953. Click **OK**.

  Note
  You cannot apply EIGRP filtering rules on VTI interfaces.

- (This option is enabled only if you choose **Outbound** filter direction.) To select the protocol to which the filtering rule applies, click **Protocol** and then select BGP, RIP, Static, Connected, or OSPF from the **Select Protocol** drop-down list. You can specify the relevant **Process ID** for BGP and OSPF protocols. Select the access list from the **Access List** drop-down list. The list defines the networks that are to be received and suppressed in routing updates. To add a new standard access list object, click +. For more information, refer to Configure a standard ACL object, on page 2953. Click **OK**.

Step 5 Click **Save**.

The EIGRP filter rule is configured and applied to control route filtering based on the specified direction, interface, protocol, and access list criteria.

Configure EIGRP redistribution settings

Define rules to redistribute routes from other routing protocols to the EIGRP routing process to enable route sharing across different routing protocols.

To configure EIGRP redistribution settings, perform these steps:

Procedure

Step 1 On the **EIGRP** page, click the **Redistribution** tab.

Step 2 Click **Add**.

Step 3 In the **Add Redistribution** dialog box, from the **Protocol** drop-down list, select the source protocol from which the routes are being redistributed:

- **BGP**: Redistributes routes discovered by the BGP routing process to EIGRP. Specify the relevant **Process ID**.
- **RIP**: Redistributes routes discovered by the RIP routing process to EIGRP.
- **Static**: Redistributes static routes to the EIGRP routing process. Static routes within a network statement scope automatically redistribute to EIGRP, without additional rules. Specify the metric when redistributing static routes to VTI interfaces in EIGRP. You do not need to specify the metric for static routes pointing to other types of interfaces.
- **Connected**: Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed to EIGRP; you do not need to define a redistribution rule for them.
- **OSPF**: Redistributes routes discovered by the OSPF routing process to EIGRP. Specify the relevant **Process ID**. If you select this protocol, the Match options on this dialog box become available under **Optional OSPF Redistribution**:
  - **Internal**: Routes that are internal to a specific AS.
  - **External1**: Routes that are external to the AS and imported into OSPF as a Type 1 external route.
  - **External2**: Routes that are external to the AS and imported into the selected process as a Type 2 external route.
  - **Nsaa-External1**: Not-So-Stubby Area (NSSA) routes that are external to the AS and imported into the selected process as Type 1 external routes.
  - **Nsaa-External2**: (NSSA) routes that are external to the AS and imported into the selected process as Type 2 external routes.

Note
These options are not available when redistributing static, connected, RIP, or BGP routes.

Step 4 Under **Optional Metrics**, enter the relevant values:

- **Bandwidth**: The minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295.

- **Delay Time**: The routing delay in tens of microseconds. Valid values range from 0 to 4294967295.
- **Reliability**: The likelihood of successful packet transmission is expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability.
- **Loading**: The effective bandwidth of the route. Valid values range from 1 to 255. 255 indicates 100 percent loading.
- **MTU**: The smallest permissible value for the maximum transmission unit of the path. Valid values range from 1 to 65535.
- **Route Map**: Select the route map object to apply to the redistribution entry from the **Route Map** drop-down list. To create a new route map object, click +. For instructions on adding a new route map, refer to Configure Route Map Entry.

Step 5 Click **OK**.

Step 6 Click **Save**.

The EIGRP redistribution settings have been configured and saved. Routes from the selected source protocol will be redistributed to EIGRP according to the specified rules and metrics.

Configure EIGRP summary address settings

You can configure summary addresses for each interface. If you want to create summary addresses that do not occur at a network boundary, you should manually define them. Alternatively, you can use summary addresses with automatic route summarization disabled. If more specific routes are available in the routing table, EIGRP advertises the summary address with a metric equal to the minimum of all specific routes.

Procedure

Step 1 On the **EIGRP** page, click the **Summary Address** tab.

Step 2 Click +**Add**.

Step 3 From the **Interface** drop-down list, select the interface from which the summary address is advertised.

Step 4 From the **Network** drop-down list, select the network object with a specific IP address and network mask to be summarized. To add a new network, click +. For detailed instructions for adding networks, refer to Networks, on page 2982.

Step 5 In the **Administrative Distance** field, enter the administrative distance of the summary route. Valid values range from 1 to 255.

Step 6 Click **OK**.

Step 7 Click **Save**.

The EIGRP summary address settings are configured and saved. EIGRP will advertise the summary address from the specified interface with the configured administrative distance.

Configure EIGRP interfaces settings

Manage interface settings to optimize routing operations.

To configure EIGRP interface settings, perform these steps:

Procedure

Step 1 On the **EIGRP** page, click the **Interfaces** tab.

Step 2 Click **Add**.

Step 3 Select the appropriate interface from the **Interface** drop-down list.

Step 4 In the **Hello Interval** field, enter the hello interval in seconds for EIGRP packets on an interface. Valid values range from 1 to 65535. The default value is 5 seconds.

Step 5 In the **Hold Time** field, enter the hold time that is advertised by the device in EIGRP hello packets. Valid values range from 1 to 65535. The default value is 15 seconds.

To enable EIGRP split-horizon on the interface, check the **Split Horizon** check box.

Step 6 In the **Delay Time** field, enter the delay time in tens of microseconds. Valid values are from 1 to 16777215.

Note
This option is not supported in multi-context mode devices.

Step 7 Specify authentication properties:

- **Enable MD5 Authentication**: Check the check box to use the MD5 hash algorithm for authentication of EIGRP packets.
- **Key Type**: Select one key type from the drop-down list:
  - **None**: Indicates that no authentication is required.
  - **Unencrypted**: Indicates that the key string is a clear text password for authentication.
  - **Encrypted**: Indicates that the key string is an encrypted password for authentication.
  - **Auth Key**: Indicates that the key string is an EIGRP authentication key.
- **Key ID**: The ID of the key that is used to authenticate EIGRP updates. Enter a numerical key identifier. Valid values range from 0 to 255.
- **Key**: An alphanumeric character string up to 17 characters. For an encrypted authentication type, this field should have a minimum of 17 characters.
- **Confirm Key**: Re-enter the key.

Step 8 Click **OK**.

Step 9 Click **Save**.

The EIGRP interface settings that you configured are saved. The interface uses the specified hello interval, hold time, split-horizon setting, delay time, and authentication parameters for EIGRP routing operations.

Configure EIGRP advanced settings

Configure EIGRP advanced settings to optimize routing behavior and performance in your network.

To configure EIGRP advanced settings, perform these steps:

Procedure

Step 1 On the **EIGRP** page, click the **Advanced** tab.

Step 2 Under **Default Route Information**, specify the sending and receiving of default route information in EIGRP updates.

- (Appears for non-cluster and cluster in spanned etherchannel mode)**Router ID (IP Address)**: Enter the ID used to identify the originating router for external routes. If an external route is received with the local router ID, the route is discarded. Specify a global address for the router ID to prevent this issue. Each EIGRP router requires a unique value.
- (Appears only for a cluster in individual interface mode)**IPv4 Address Pool**: Select the relevant cluster pool value (IPv4 address pool object). To create the address pool, refer to Configure address pools, on page 2954.
- **Accept Default Route Info**: Check the check box to configure EIGRP to accept default route information.
  - **Access List**: From the **Access List** drop-down list, specify a standard access list that defines the networks that are allowed and the networks that are not when receiving default route information.

    To add a new standard access list object, click +. For more information, refer to Configure a standard ACL object, on page 2953.
- **Send Default Route Info**: Check the check box to configure EIGRP to advertise exterior default routing information.
  - **Access List**: From the **Access List** drop-down list, specify a standard access list that defines the networks that are allowed and the networks that are not when sending default route information.

    To add a new standard access list object, click +. For more information, refer to Configure a standard ACL object, on page 2953.

Step 3 Under **Administrative Distance**, specify:

- **Internal Distance**: Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system. Valid values range from 1 to 255. The default value is 90.
- **External Distance**: Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a neighbor external to the autonomous system. Valid values range from 1 to 255. The default value is 170.

Step 4 Under **Adjacency Changes**, specify:

- **Log Neighbor Changes**: Check the check box to enable the logging of EIGRP neighbor adjacency changes.

- **Log Neighbor Warnings**: Check the check box to enable the logging of EIGRP neighbor warning messages.
- (Optional) Enter the time interval (in seconds) between repeated neighbor warning messages. Valid values range from 1 to 65535. Repeated warnings are not logged if they occur during this interval.

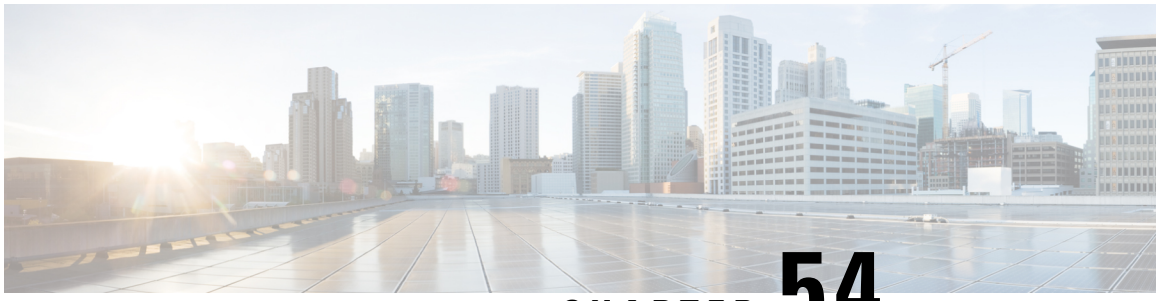
Step 5 Under **Stub**, to enable the device as an EIGRP stub routing process, click one or more of the EIGRP stub routing processes check boxes:

- **Receive only**: Configures the EIGRP stub routing process to receive route information from the neighbor routers but not send route information to the neighbors. If this option is selected, you cannot select any of the other stub routing options.
- **Connected**: Advertises connected routes.
- **Redistributed**: Advertises redistributed routes.
- **Static**: Advertises static routes.
- **Summary**: Advertises summary routes.

Step 6 Under **Default Metrics**, define the default metrics for routes redistributed to the EIGRP routing process:

- **Bandwidth**: The minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295.
- **Delay Time**: The route delay in tens of microseconds. Valid values range from 0 to 4294967295.
- **Reliability**: The likelihood of successful packet transmission expressed as a number from 0 to 255. The value 255 indicates 100 percent reliability; 0 means no reliability.
- **Loading**: The effective bandwidth of the route. Valid values range from 1 to 255; 255 indicates 100 percent loading.
- **MTU**: The smallest allowed value for the maximum transmission unit of the path. Valid values range from 1 to 65535.

EIGRP advanced settings are configured according to your specifications, optimizing routing behavior and performance.



CHAPTER 54

BGP

This section describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information using the Border Gateway Protocol (BGP).

- BGP, on page 1491
- Requirements and prerequisites for BGP, on page 1495
- Guidelines for BGP, on page 1495
- Configure BGP, on page 1496

BGP

Border Gateway Protocol (BGP) is an inter- and intra-autonomous system routing protocol that

- operates within and between autonomous systems (networks or groups of networks under common administration with common routing policies),
- exchanges routing information for the Internet, and
- serves as the protocol used between Internet service providers (ISPs).

Routing tables

A routing table is a network database that

- records paths to destination networks based on exchanged routing information
- updates dynamically as neighbors share new or modified routing routes, and
- determines the optimal path to a destination without requiring periodic updates.

BGP attributes for route selection

Routes learned via BGP have several attributes that determine the best path when multiple routes to a destination exist. The most important BGP attributes include:

- Weight: A Cisco-defined attribute that is local to a router. It is not advertised to neighboring routers. Among multiple routes to the same destination, the route with the highest weight is preferred.

- Local preference: This attribute selects an exit point from the local autonomous system (AS) and is shared throughout the local AS. If multiple exit points exist, the one with the highest local preference attribute is chosen for a specific route.
- Multi-exit discriminator: The multi-exit discriminator (MED) or metric attribute is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. It is referred to as a suggestion because the external AS that is receiving the MEDs may also be using other BGP attributes for route selection. The route with the lower MED metric is preferred.
- Origin: The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three possible values and is used in route selection.
  - IGP: The route is interior to the originating AS. This value is set when the network router configuration command is used to inject the route into BGP.
  - EGP: The route is learned via the Exterior Border Gateway Protocol (EBGP).
  - Incomplete: The origin of the route is unknown or learned in some other way. An origin of incomplete occurs when a route is redistributed into BGP.
- AS_PATH: When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. Only the route with the shortest AS_PATH list is installed in the IP routing table.
- Next hop: The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers. For IBGP, the EBGP next-hop address is carried into the local AS. However, when the next hop is in the same subnet the peering address of the EBGP peer, the next hop is not modified. This behavior is referred to as the third-party next hop.

  Use the **next-hop-self** command when redistributing VPN-advertised routes to IBGP as peers to ensure that the routes are redistributed with the correct next hop IP.
- Community: The community attribute provides a way of grouping destinations, called communities, to which routing decisions (such as acceptance, preference, and redistribution) can be applied. Route maps are used to set the community attribute. The predefined community attributes are as follows:
  - no-export: Do not advertise this route to EBGP peers.
  - no-advertise: Do not advertise this route to any peer.
  - internet: Advertise this route to the Internet community; all routers in the network belong to it.



Note As loop detection is done by scanning the full AS_PATH (as specified in the AS_PATH attribute), and checking that the AS number of the local system does not appear in the AS_PATH. By default, EBGP advertises the learned routes to the same peer to prevent additional CPU cycles on the device in performing loop checks and to avoid delays in the existing outgoing update tasks.

BGP usage scenarios

BGP usage scenarios are network configurations that

- enable ISPs to exchange customer and ISP routes between autonomous systems using External BGP (EBGP),
- allow service providers to exchange routes within an autonomous system using Interior BGP (IBGP), and
- support IPv6 prefix routing over IPv6 networks.

BGP implementation contexts

Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP) such as OSPF for the exchange of routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

BGP can also be used for carrying routing information for IPv6 prefix over IPv6 networks.

BGP path selection

BGP path selection is a routing process that selects only one path as the best path when BGP receives multiple advertisements for the same route from different sources, puts the selected path in the IP routing table, and propagates the path to its neighbors using specific criteria in order of precedence.

BGP path selection criteria

BGP uses these criteria, in the order presented, to select a path for a destination:

1. If the path specifies a next hop that is inaccessible, drop the update. Otherwise, prefer the path with the largest weight.
2. If the weights are the same, prefer the path with the largest local preference. If the local preferences are the same, prefer the path that was originated by BGP running on this router.
3. If no route was originated, prefer the route that has the shortest AS_PATH. If all paths have the same AS_PATH length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
4. If the origin codes are the same, prefer the path with the lowest MED attribute. If the paths have the same MED, prefer the external path over the internal path.
5. If the paths are still the same, prefer the path through the closest IGP neighbor.
6. Determine if multiple paths require installation in the routing table for BGP multipath, on page 1493.
7. If both paths are external, prefer the path that was received first (the oldest one).
8. Prefer the path with the lowest IP address, as specified by the BGP router ID. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length.
9. Prefer the path that comes from the lowest neighbor address.

BGP multipath

BGP multipath is a BGP feature that

- allows installation into the IP routing table of multiple equal-cost BGP paths to the same destination prefix,
- enables traffic to the destination prefix to be shared across all installed paths, and
- does not affect best path selection algorithm.

BGP multipath requirements

These paths are installed in the table together with the best path for load-sharing. BGP multipath does not affect best path selection. For example, a router still designates one of the paths as the best path, according to the algorithm, and advertises this best path to its BGP peers.

In order to be candidates for multipath, paths to the same destination need to have these characteristics equal to the best-path characteristics:

- Weight
- Local preference
- AS-PATH length
- Origin code
- Multi Exit Discriminator (MED)
- One of these:
  - Neighboring AS or sub-AS (before the addition of the BGP Multipaths)
  - AS-PATH (after the addition of the BGP multipaths)

Some BGP multipath features put additional requirements on multipath candidates:

- The path should be learned from an external or confederation-external neighbor (eBGP).
- The IGP metric to the BGP next hop should be equal to the best-path IGP metric.

These are the additional requirements for internal BGP (iBGP) multipath candidates:

- The path should be learned from an internal neighbor (iBGP).
- The IGP metric to the BGP next hop should be equal to the best-path IGP metric, unless the router is configured for unequal-cost iBGP multipath.

BGP inserts up to *n* most recently received paths from multipath candidates into the IP routing table, where *n* is the number of routes to install in the routing table, as specified when you configure BGP multipath. The default value, when multipath is disabled, is 1.

For unequal-cost load balancing, you can also use BGP Link Bandwidth.

Note The equivalent next-hop-self is performed on the best path that is selected among eBGP multipaths before it is forwarded to internal peers.

Requirements and prerequisites for BGP

This reference provides the requirements and prerequisites necessary for configuring BGP on supported devices.

Model support

Firewall Threat Defense

Firewall Threat Defense Virtual

Supported domains

Any

User roles

Admin

Network Admin

Guidelines for BGP

Context mode guidelines

Consider these context mode requirements when implementing BGP.

- Supported in single and multiple context modes.
- Only one Autonomous System (AS) number is supported for all contexts.

Firewall mode guidelines

Does not support transparent firewall mode. BGP is supported only in routed mode.

IPv6 guidelines

Supports IPv6.

Additional guidelines

Consider these additional operational requirements and limitations.

- For BGP, the next hop IP address for the routes is the network IP address and not 0.0.0.0.
- The system does not add a route entry for the IP address received over PPPoE in the CP route table. BGP always looks into the CP route table for initiating the TCP session, hence BGP does not form TCP session.

  Thus, BGP over PPPoE is not supported.
- BGP is not supported on management-only or BVI interfaces.
- To avoid adjacency flaps due to route updates being dropped if the route update is larger than the minimum MTU on the link, ensure that you configure the same MTU on the interfaces on both sides of the link.

- BGP with PATH MTU (PMTU) can cause adjacency flaps if MTU discovery fails, especially with ECMP routing. Hence, be cautious while using BGP, PMTU, and ECMP as packet drops can occur if MTU discovery fails due to any reason.
- The BGP table of the member unit is not synchronized with the control unit table. Only its routing table is synchronized with the control unit routing table.
- Enable the graceful timeout feature when you use BGP on Firewall Threat Defense devices that are configured in high availability mode. Enabling this feature reduces packet drops that occur due to issues related to BGP convergence.
- When you configure a route-based site-to-site VPN using static or dynamic VTI interfaces, ensure that the TTL hop count is greater than one if you use BGP as the routing protocol.
- On Firewall Threat Defense devices, changing a large number of BGP neighbors from iBGP to eBGP in a single deployment can fail. During deployment, Cloud-Delivered Firewall Management Center can remove and re-create the same peer in immediate succession, and the deployment can fail if the previous peer deletion is still in progress. To avoid this issue, split the changes across multiple deployments. Remove or modify a subset of neighbors, deploy the configuration, and then apply additional neighbor changes in subsequent deployments.

Configure BGP

Configure BGP routing protocol to establish peering relationships and control route advertisement between autonomous systems.

BGP configuration involves multiple components including basic settings, neighbor relationships, route filtering, and redistribution policies.

Procedure

Step 1 Configure basic BGP settings, on page 1497

Step 2 Configure BGP general settings, on page 1499

Step 3 Configure BGP Neighbor Settings, on page 1500

Step 4 Configure BGP aggregate address settings, on page 1504

Step 5 Configure BGPv4 filtering settings, on page 1505

Note
The Filtering section is applicable only to IPv4 settings

Step 6 Configure BGP network settings, on page 1506

Step 7 Configure BGP redistribution settings, on page 1507

Step 8 Configure BGP route injection settings, on page 1508

Step 9 Configure BGP route import and export, on page 1509

BGP is configured with all necessary settings, enabling dynamic routing and proper route exchange between autonomous systems.

Configure basic BGP settings

Enable and customize BGP settings to control routing behavior and improve network operations.

You can set many basic settings for BGP.

For a device using virtual routing, the basic settings described in this section must be configured in the **BGP** page under **General Settings**. For more information, refer to Modifications to the Cloud-Delivered Firewall Management Center web interface routing page.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** and do the following:

a) (For a virtual-router-aware device) Under **General Settings**, click **BGP**.

b) Check the **Enable BGP** check box to enable the BGP routing process.

c) In the **AS Number** field, enter the autonomous system (AS) number for the local BGP process. The AS number internally includes multiple autonomous numbers. Valid values are from 1 to 4294967295 or from 1.0 to 65535.65535 when a dot notation is used. The AS number is a uniquely assigned value that identifies each network on the Internet.

d) In the **Router Id** drop-down list, choose **Automatic** or **Manual** (appears for non-cluster and a cluster in spanned etherchannel mode) or **Cluster Pool** (appears for a cluster in individual interface mode). If you choose **Automatic**, the highest-level IP address on the Firewall Threat Defense device is used as the router ID. If you choose **Manual**, enter the IP address in the **IP Address** field. If you choose **Cluster Pool**, enter the cluster pool value in the **Cluster Pool** field. For more information on creating the cluster pool address, refer to Configure address pools, on page 2954..

e) To use a fixed router ID, choose **Manual** and enter an IPv4 address in the **IP Address** field. The default value is **Automatic**. For a virtual router-aware device, you can override the router ID settings in the **Virtual Routers** > **BGP** page.

Step 3 (Optional) Edit the various BGP settings, starting with **General**. The defaults for these settings are appropriate in most cases, but you can adjust them to fit the needs of your network. Click **Edit** ( ) to edit the settings in the group:

a) Enter a **Scanning Interval** for BGP routers for next-hop validation. Valid values are from 5 to 60 seconds. The default value is 60.

b) In the **Number of AS numbers in AS_PATH attribute of received routes** field, enter the maximum number of AS numbers allowed in the AS_PATH attribute of received routes. An AS_PATH attribute is a sequence of intermediate AS numbers between source and destination routers that form a directed route for packets to travel. Valid values are between 1 and 254. The default value is **None**.

c) Check the **Log Neighbor Changes** check box to enable logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability. This is enabled by default.

d) Check the **Use TCP path MTU discovery** check box to use the Path MTU determining technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation. This is enabled by default.

e) Check the **Reset session upon failover** check box to reset the external BGP session immediately upon link failure. This is enabled by default.

f) Check the **Enforce the first AS is peer's AS for EBGP routes** check box to discard incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH

attribute. This prevents a misconfigured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system. This is enabled by default.

g) Check the **Use dot notation for AS number** check box to split the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 1 to 65535 are represented as decimal numbers, and AS numbers greater than 65535 are represented using dot notation. This is disabled by default.

h) Enter a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60. The default value is 30.

i) Click **OK**.

Step 4 (Optional) Edit the **Best Path Selection** section:

a) Enter a value for **Default local preference** between 0 and 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers and access servers in the local autonomous system.

b) Check the **Allow comparing MED from different neighbors** check box to allow the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. This is disabled by default.

c) Check the **Compare Router ID for identical EBGP paths** check box to compare similar paths received from external BGP peers during the best path selection process and switch the best path to the route with the lowest router ID. This is disabled by default.

d) Check the **Pick the best-MED path among paths advertised by neighbor AS** check box to enable MED comparison among paths learned from confederation peers. The comparison between MEDs is made only if no external autonomous systems are there in the path. This is disabled by default.

e) Check the **Treat missing MED as the least preferred path** check box to consider the missing MED attribute as having a value of infinity, making the path the least desirable; therefore, a path with a missing MED is least preferred. This is disabled by default.

f) Click **OK**.

Step 5 (Optional) Edit the **Neighbor Timers** section:

a) Enter the time interval for which the BGP neighbor remains active after not sending a keepalive message in the **Keepalive interval** field. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.

b) Enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured in the **Hold time** field. The default value is 180 seconds. Specify a value from 0 to 65535.

c) (Optional) Enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured in the **Min Hold time** field. Specify a value from 3 to 65535.

Note
A hold time of less than 20 seconds increases the possibility of peer flapping.

d) Click **OK**.

Step 6 In the **Next Hop** section, optionally check the **Address tracking** check box to enable BGP next hop address tracking and enter the **Delay Interval** between checks on updated next-hop routes installed in the routing table. Click **OK**.

Note
The **Next Hop** section is applicable only to IPv4 settings.

Step 7 (Optional) Edit the **Graceful Restart** section:

Note
This section is available only when the Firewall Threat Defense device is in failover or spanned cluster mode. This is done so that there is no drop in packets in the traffic flow, when one of the devices in the failover setup fails.

a) Check the **Enable Graceful Restart** check box to enable Firewall Threat Defense peers to avoid a routing flap following a switchover.
b) Specify the time duration that Firewall Threat Defense peers will wait to delete stale routes before a BGP open message is received in the **Restart Time** field. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
c) Enter the time duration that the Firewall Threat Defense will wait before deleting stale routes after an end of record (EOR) message is received from the restarting Firewall Threat Defense in the **Stalepath Time** field. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.
d) Click **OK** and then click **Save**.

Step 8 To view the BGP basic settings, from the virtual routers drop-down, select the desired router, and then click **BGP**.

This page displays the basic settings that are configured in the **Settings** page. You can edit the router ID settings on this page.

Step 9 To edit the router ID settings, modify the IP address in the **IP Address** fields. The modified value overrides the router ID settings that were configured on the **BGP** page under **General Settings**.

The device is configured with basic BGP settings, enabling BGP routing, specifying the AS number, router ID, and any custom parameters you set.

Configure BGP general settings

Configure route maps, administrative route distances, synchronization, next-hop, and packet forwarding for BGP to optimize routing performance in your network.

The defaults for these settings are appropriate in most cases, but you can adjust them to fit the needs of your network.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, select the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 In the **General** tab, update the following sections:

a) In the **Setting** section, enter or select a **RouteMap** object and click **OK**.

Note
The **RouteMap** field is applicable only to IPv4 settings.

b) In the **Administrative Route Distances** section, update the following as required, and click **OK**:

- **External**: Enter the administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255. The default value is 20.
- **Internal**: Enter administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255. The default value is 200.
- **Local**: Enter administrative distance for local BGP routes. Local routes are routes originated by the local BGP process, such as networks specified with the network router configuration command or routes redistributed into BGP from another routing process. The range of values for this argument is from 1 to 255. The default value is 200.

c) In the **Routes and Synchronization** section, update the following as required, and click **OK**:

- (Optional) **Generate default routes**: Check the check box of this option to configure default-information originate.
- (Optional) **Summarize subnet routes into network level routes**: Check the check box of this to configure automatic summarization of subnet routes into network-level routes. This check box is applicable only to IPv4 settings.
- (Optional) **Advertise inactive routes**: Check the check box of this to advertise routes that are not installed in the routing information base (RIB).
- (Optional) **Synchronize between BGP and IGP systems**: Check the check box of this to enable synchronization between BGP and your Interior Gateway Protocol (IGP) system. Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems.
- (Optional) **Redistribute IBGP into IGP**: Check the check box of this to configure IBGP redistribution into an interior gateway protocol (IGP), such as OSPF.

d) In the **Forward Packets Over Multiple Paths** section, update the following as required and click **OK**:

- (Optional) **Number of Paths**: Enter the maximum number of Border Gateway Protocol routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.
- (Optional) **IBGP Number of Paths**: Enter the maximum number of parallel internal Border Gateway Protocol (IBGP) routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.

Step 5 Click **Save**.

BGP general settings are configured with your specified route maps, administrative distances, synchronization options, and multi-path forwarding parameters.

Configure BGP Neighbor Settings

Define and customize BGP neighbor connections on your device for efficient routing and advanced peer management.

A BGP router must connect with each of its peers before exchanging updates. These peers are called BGP neighbors. Perform this procedure to define BGP IPv4 or IPv6 neighbors and neighbor settings.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

(For a virtual router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring BGP.

Step 2 Choose **BGP** > **IPv4** or **IPv6**. Click **Neighbor**.

Step 3 Click **Add** to define BGP neighbors and neighbor settings.

a) Enter the BGP neighbor **IP address**. This IP address is added to the BGP neighbor table. When you are configuring BGP IPv6 on static VTI, enter the virtual tunnel IP address of the neighbor. Choose the BGP neighbor **Interface**.

Note
The **Interface** field is only applicable to IPv6 settings.

b) Enter the autonomous system to which the BGP neighbor belongs, in the **Remote AS** field. Check the **Enabled address** check box to enable communication with this BGP neighbor. Further neighbor settings will be configured only if the **Enabled address** check box is selected.

c) (Use this option only when the routing does not result in a loop.) Check the **AS Override** check box to enable overriding of the AS number of the originating router with the AS number of the sending BGP router.

Note
Loop prevention in BGP is achieved by verifying the AS number in the AS path. BGP rejects route updates when the AS path attribute contains its own AS number. However, in circumstances where the same ASN is used by discontiguous network segments, to establish end-to-end reachability, you can opt to override the AS number with that of the sending BGP router.

d) (Optional) Check the **Shutdown administratively** check box to disable a neighbor or peer group.

e) (Optional) Check the **Configure graceful restart (failover / spanned mode)** check box to enable configuration of the BGP graceful restart capability for this neighbor. After selecting this option, you must check the **Enable graceful restart** check box to specify whether graceful restart should be enabled or disabled for this neighbor.

Note
- The graceful restart is enabled only when the device is in HA mode or when L2 cluster (all nodes from the same network) is configured.
- The graceful restart option for BGPv6 is enabled only on Firewall Threat Defense Version 7.3+.
- If you configure graceful restart only at General Settings and not at BGP IPv6, the global General Settings configuration persists.
- If you configure graceful restart at General Settings and also at BGP IPv6, the global General Settings configuration is overridden by the BGP IPv6 configuration settings.

f) (Optional) To enable configuration of the BFD support for BGP, from the **BFD Fallover** drop-down list, choose the BFD type—single-hop, multi-hop, or auto-detect-hop. This selection registers the BGP neighbor

to receive forwarding path detection failure messages from BFD. Choose None if you do not want to have BFD support.

g) (Optional) Enter a **Description** for the BGP neighbor.

h) (Optional) From the **Update Source** drop-down list, choose an interface to source the BGP packets.

You can choose a loopback address as this interface to overcome path failures. You can also choose any physical, port-channel, or a sub-interface.

Step 4 (Optional) In **Filtering Routes**, use access lists, route maps, prefix lists and AS path filters as required, to distribute BGP Neighbor information. Update the following sections:

a) Choose or select the appropriate incoming or outgoing **Access List** to distribute BGP neighbor information.

Note
Access lists are only applicable to IPv4 settings.

b) Choose or select the appropriate incoming or outgoing **Route Maps** to apply a route map to incoming or outgoing routes.

c) Choose or select the appropriate incoming or outgoing **Prefix List** to distribute BGP neighbor information.

d) Choose or select the appropriate incoming or outgoing **AS path filter** to distribute BGP neighbor information.

e) Check the check box of **Limit the number of prefixes allowed from the neighbor** to control the number of prefixes that can be received from a neighbor.

- Enter the maximum number of prefixes allowed from a specific neighbor in the **Maximum Prefixes** field.
- Enter the percentage (of maximum) at which the router starts to generate a warning message in the **Threshold Level** field. Valid values are integers between 1 and 100. The default value is 75.

f) Check the **Control prefixes received from the peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following:

- Check the **Terminate peering when prefix limit is exceeded** check box to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the **Restart interval** field.
- Check the **Give only warning message when prefix limit is exceeded** check box to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.

g) Click **OK**.

Step 5 (Optional) In **Routes**, specify miscellaneous Neighbor route parameter. Proceed to update the following:

a) Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisement Interval** field. Valid values are between 1 and 600.

b) Check the **Remove private AS numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.

c) Check the **Generate default routes** check box to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route. Enter or Select the route map that allows the route 0.0.0.0 to be injected conditionally in the **Route map** field.

d) To add conditionally advertised routes, click Add Row +. In the Add Advertised Route dialog box, do the following:

1. Add or choose a route map in the **Advertise Map** field, that will be advertised if the conditions of the exist map or the non-exist map are met.
2. Click **Exist Map** and choose a route map from the Route Map Object Selector. This route map is compared with the routes in the BGP table, to determine whether the advertise map route is advertised.
3. Click **Non-Exist Map** and choose a route map from the Route Map Object Selector. This route map is compared with the routes in the BGP table, to determine whether the advertise map route is advertised.
4. Click **OK**.

Step 6 In **Timers**, check the **Set timers for the BGP peer** check box to set the keepalive frequency, hold time, and minimum hold time.

- **Keepalive interval**: Enter the frequency (in seconds) with which Firewall Threat Defense sends keepalive messages to the neighbor. Valid values are between 0 and 65535. The default value is 60 seconds.
- **Hold time**: Enter the interval (in seconds) after not receiving a keepalive message that Firewall Threat Defense declares a peer dead. Valid values are between 0 and 65535. The default value is 180 seconds.
- **Min hold time**: (Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that Firewall Threat Defense declares a peer dead. Valid values are between 3 and 65535. The default value is 3 seconds.

  Note
  A hold time of less than 20 seconds increases the possibility of peer flapping.

Step 7 In **Advanced**, update the following:

a) (Optional) Check the **Enable Authentication** check box to enable MD5 authentication on a TCP connection between two BGP peers.

1. Choose an encryption type from the **Enable Encryption** drop-down list.
2. Enter a password in the **Password** field. Reenter the password in the **Confirm Password** field. The password is case-sensitive and can be up to 25 characters long when the service password-encryption command is enabled and up to 81 characters long when the service password-encryption command is not enabled. The string can contain any alphanumeric characters, including spaces.

   Note
   You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

b) (Optional) Select the **Send Community attribute to this neighbor** check box to specify that community attributes should be sent to the BGP neighbor

c) (Optional) Select the **Use FTD as next hop for this neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.

d) Select the **Disable Connection Verification** check box to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IP address. When deselected (default), a BGP routing process will verify the connection of single-hop eBGP peering session (TTL=254) to determine if the eBGP peer is directly connected to the same network segment by default. If the peer is not directly

connected to same network segment, connection verification will prevent the peering session from being established.

e) Select **Allow connections with neighbor that is not directly connected** to accept and attempt BGP connections to external peers residing on networks that are not directly connected. (Optional) Enter the time-to-live in the **TTL hops** field. Valid values are between 1 and 255. However, for eBGP multi-hop neighborship configuration, you must enter the **TTL hops** value. Alternatively, select **Limited number of TTL hops to neighbor**, to secure a BGP peering session. Enter the maximum number of hops that separate eBGP peers in the **TTL hops** field. Valid values are between 1 and 254.

f) (Optional) Select the **Use TCP MTU path discovery** check box to enable a TCP transport session for a BGP session.

g) Choose the TCP connection mode from the **TCP Transport Mode** drop-down list. Options are Default, Active, or Passive.

h) (Optional) Enter a **Weight** for the BGP neighbor connection.

i) Select the **BGP Version** that Firewall Threat Defense will accept from the drop-down list. The version can be set to 4-Only to force the software to use only Version 4 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

Step 8 Update **Migration**, only if AS migration is considered.

Note
The AS migration customization should be removed after transition has been completed.

a) (Optional) Check the **Customize the AS number for routes received from the neighbor** check box to customize the AS_PATH attribute for routes received from an eBGP neighbor.

b) Enter the local autonomous system number in the **Local AS number** field. Valid values are any valid autonomous system number from 1 to 4294967295 or 1.0 to 65535.65535.

c) (Optional) Check the **Do not prepend local AS number to routes received from neighbor** check box to prevent the local AS number from being prepended to any routes received from eBGP peer.

d) (Optional) Check the **Replace real AS number with local AS number in routes received from neighbor** check box to replace the real autonomous system number with the local autonomous system number in the eBGP updates. The autonomous system number from the local BGP routing process is not prepended.

e) (Optional) Check the **Accept either real AS number or local AS number in routes received from neighbor** check box to configure the eBGP neighbor to establish a peering session using the real autonomous system number (from the local BGP routing process) or by using the local autonomous system number.

Step 9 Click **OK**, and then click **Save**.

The BGP neighbor and its associated advanced settings are configured. The device can now establish and manage BGP sessions as specified, ensuring effective routing and network reachability.

Configure BGP aggregate address settings

Route aggregation combines the attributes of several different routes so that only a single route is advertised, reducing the number of routes that need to be advertised.

BGP neighbors store and exchange routing information and the amount of routing information increases as more BGP speakers are configured. Route aggregation is the process of combining the attributes of several different routes so that only a single route is advertised. Aggregate prefixes use the classless interdomain routing (CIDR) principle to combine contiguous networks into one classless SET of IP addresses that can be

summarized in routing tables. As a result fewer routes need to be advertised. Use the Add/Edit Aggregate Address dialog box to define the aggregation of specific routes into one route.

Procedure

Step 1 When editing the Firewall Threat Defense device, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Add Aggregate Address**.

Step 5 Enter a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60. The default value is 30.

Step 6 Click (+)**Add** and update the **Add Aggregate Address** dialog box:

a) **Network**: Enter an IPv4 address or select the desired network/hosts objects.
b) **Attribute Map**: (Optional) Enter or select the route map used to SET the attribute of the aggregate route.
c) **Advertise Map**: (Optional) Enter or select the route map used to select the routes to create AS_SET origin communities.
d) **Suppress Map**: (Optional) Enter or select the route map used to select the routes to be suppressed.
e) **Generate AS SET path Information**: (Optional) Check the check box to enable generation of autonomous system SET path information.
f) **Filter all routes from updates**: (Optional) Check the check box to filter all more-specific routes from updates.
g) Click **OK**.

The BGP aggregate address settings are configured and fewer routes need to be advertised due to route aggregation.

What to do next

- For BGPv4 settings, proceed to Configure BGPv4 filtering settings, on page 1505.
- For BGPv6 settings, proceed to Configure BGP network settings, on page 1506.

Configure BGPv4 filtering settings

Configure filtering settings to filter routes or networks received in incoming BGP updates and restrict routing information that the router learns or advertises.

Filtering settings are used to filter routes or networks received in incoming BGP updates. Filtering is used to restrict routing information that the router learns or advertises.

Before you begin

Filtering is only applicable for a BGP IPv4 routing policy.

Follow these steps to configure BGPv4 filtering settings:

Procedure

Step 1 On the Device Management page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4**.

Step 4 Click **Filtering**.

Note
The **Filtering** field is applicable only to IPv4 settings.

Step 5 Click (+)**Add** and update the **Add Filter** dialog box:

a) **Access List**: Choose an access control list that defines which networks are to be received and which are to be suppressed in routing updates.

b) **Direction**: (Optional) Choose a direction that specifies if the filter should be applied to inbound updates or outbound updates.

c) **Protocol**: (Optional) Choose the routing process for which you want to filter: None, BGP, Connected, OSPF, RIP, or Static.

d) **Process ID**: (Optional) Enter the process ID for the OSPF routing protocol.

e) Click **OK**.

Step 6 Click **Save**.

BGPv4 filtering settings are configured to filter routes or networks based on the specified access control list and filtering criteria.

Configure BGP network settings

Configure BGP network settings to control which networks are advertised by the BGP routing process and to apply route map filters for network advertisement.

Network settings are used to add networks that will be advertised by the BGP routing process and route maps that will be examined to filter the networks to be advertised.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Networks**.

Step 5 Click **Add** and update the **Add Networks** dialog box:

a) **Network**: Choose the network to be advertised by the BGP routing processes.

Note
For a network prefix to be advertised, a route to the device must exist on the routing table.

To add a new network object, refer to Create a network object, on page 2985.

b) (Optional) **Route Map**: Enter or choose a route map that should be examined to filter the networks to be advertised. If not specified, all networks are redistributed. To add a new route map object, refer to Configure a route map, on page 3017.

c) Click **OK**.

Step 6 Click **Save**.

The BGP network settings are configured with the specified networks and route maps for advertisement filtering.

Configure BGP redistribution settings

Configure redistribution settings to define the conditions for redistributing routes from another routing domain into BGP.

Redistribution settings allow you to define the conditions for redistributing routes from another routing domain into BGP.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Redistribution**.

Step 5 Click **Add** and update the **Add Redistribution** dialog:

a) **Source Protocol**: Select the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.

Note
User-defined virtual routers do not support redistributing routes from RIP.

b) **Process ID**: Enter the identifier for the selected source protocol. Applies to the OSPF protocol. For devices using virtual routing, this drop-down lists the process ID assigned for the virtual router for which you are configuring the BGP settings.

c) **Metric**: (Optional) Enter a metric for the redistributed route.

d) **Route Map**: Enter or select a route map that should be examined to filter the networks to be redistributed.

If not specified, all networks are redistributed. To create a new route map object, click **Add** (+). For the procedure to add a new route map, see Configure Route Map Entry.

e) **Match**: The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more conditions. These options are enabled only when OSPF is chosen as the Source Protocol.

- Internal
- External 1
- External 2
- NSSA External 1
- NSSA External 2

f) Click **OK**.

The BGP redistribution settings are configured and routes from the specified source protocol are redistributed into the BGP domain according to the defined conditions.

Configure BGP route injection settings

Configure BGP route injection settings to control which routes are conditionally injected into the BGP routing table based on defined criteria.

Route injection settings allow you to define the routes to be conditionally injected into the BGP routing table.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Route Injection**.

Step 5 Click **Add** and update the **Add Route Injection** dialog box:

a) **Inject Map**: Enter or select the route map that specifies the prefixes to inject into the local BGP routing table. To create a new route map object, click **Add** (+). For the procedure to add a new route map, see Configure Route Map Entry.

b) **Exist Map**: Enter or select the route map containing the prefixes that the BGP speaker will track.

c) **Injected routes will inherit the attributes of the aggregate route**: Check this box to configure the injected route to inherit attributes of the aggregate route.

d) Click **OK**.

Step 6 Click **Save**.

The BGP route injection settings are configured and the specified routes will be conditionally injected into the BGP routing table according to the defined route maps.

Configure BGP route import and export

This task enables you to implement inter-virtual-router route leaking in BGP by importing or exporting routes using route target extended communities, allowing controlled sharing of routes between virtual routers.

In BGP, you can implement an inter-virtual-router route leak by importing or exporting routes using the route target extended community of the destination and source virtual routers respectively. You can use a route map to filter the desired route targets instead of leaking the entire routing table. You can also leak the routes of global virtual router to user-defined virtual routers and vice versa.

- You can configure BGP to leak routes between two user-defined virtual routers using the route target extended communities:
  - Tag the routes with the route targets from the source virtual router using route target export.
  - Import the routes that are matching the route targets into the destination virtual router using route target import.
  - Optionally, you can filter routes from source virtual router or to destination virtual router using export or import route maps respectively. You can configure route map with match extended community list for filtering the routes. Similarly, you can configure route map with set extended community route targets to tag the routes with the route target extended community.
- To import routes from the global virtual router to a user-defined virtual router, specify the IPv4/IPv6 route map in Global Virtual Router Import Route Map to import to the user-defined virtual router.
- To export routes from a user-defined virtual router to the global virtual router, in addition to exporting the route targets, you can also specify the Global Virtual Router Export Route Map to export from the user-defined virtual router.

The BGP inter-virtual-router route leaking supports both IPv4 and IPv6 prefixes.

Before you begin

- Create virtual routers.
- Enable BGP on the virtual router.
- Configure BGP, on page 1496.

Follow these steps to configure BGP route import/export settings:

Procedure

Step 1 On the Device Management page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 (Supported for virtual routers only) Click **Route Import/Export**.

Step 5 In the **Route Targets Import** field, enter the route target extended community that you want to match for the routes to be imported. On deployment, the routes of the destination virtual router that match this value are imported to the source virtual router's BGP table.

Note

- The route target must be in **ASN:nn** format.
- You can enter multiple route targets as comma-separated values.
- This value can range from 0:1 to 65534:65535.

Step 6 In the **Route Targets Export** field, enter the route target extended community to tag the source virtual router's routes with the route target value. On deployment, the routes of the source virtual router are tagged with this value.

Note

- The route target must be in **ASN:nn** format.
- You can enter multiple route targets as comma separated values.
- This value can range from 0:1 to 65534:65535.

Step 7 Route maps help you to narrow down the routes to be shared instead of leaking the entire routing table. Route map filtering is applied on the list of routes that are obtained with the specified route target values:

a) (Optional) Under **User Virtual Router**, choose the route map from the **Import Route Map** drop-down list to filter the routes at the destination virtual router.

Note
The user virtual router import route map is effective only when the route targets import is configured.

b) (Optional) Under **User Virtual Router**, choose the route map from the **Export Route Map** drop-down list to filter the routes at the source virtual router before the routes are exported to other virtual routers.

Note
You can use the match and set clauses in the route map with the route target extended community lists for filtering based on other criteria or tagging the routes with the route target community values. For more information, refer to Configure a route map, on page 3017.

Step 8 To share the routes between a user-defined virtual router and global virtual router, specify the route map under the **Global Virtual Router**:

a) To leak the global virtual router routes to the user-defined virtual router, select the route map from the **Import Route Map** drop-down list. The IPv4 or IPv6 route map is imported to the user-defined virtual router.

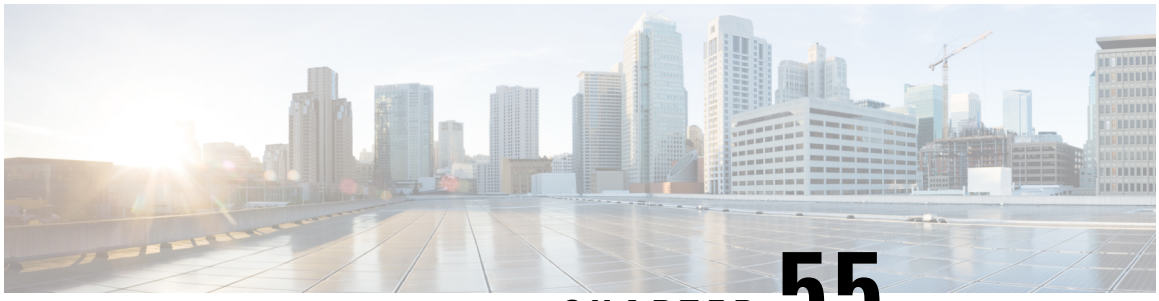
b) To leak the user-defined virtual router routes to the global virtual router, select the route map from the **Export Route Map** drop-down list. The IPv4 or IPv6 route map is exported to the global virtual router.

Note
You must specify the route targets for export apart from specifying the route map.

Note
You can use the match clause of the route map object to filter the routes for leaking. For more information, refer to Configure a route map, on page 3017.

Step 9 Follow the procedure (Step 3 to Step 8) to configure relevant BGP route import and export settings for other virtual routers as well. Click **Save**, then **Deploy** to apply your import/export settings.

When the packets flow into the ingress virtual router, BGP imports the routes from the destination virtual routers that have the matching route target value and if a route map is also configured, the routes are further filtered and used to identify the best path routes for routing the packets.



CHAPTER 55

RIP

The Routing Information Protocol (RIP) is a distance-vector routing protocol that

- enables devices to exchange routing information,
- supports authentication mechanisms, and
- facilitates redistribution of routing information.

RIP Configuration Constraints

There are constraints when deploying RIP on the Firewall Threat Defense:

- RIP configuration is restricted to the global virtual router and is not supported for user-defined virtual routers.

- RIP, on page 1513
- Requirements and prerequisites for RIP, on page 1516
- Guidelines for RIP, on page 1516
- Configure RIP, on page 1517

RIP

Routing Information Protocol (RIP) is a distance-vector routing protocol that

- uses hop count as the metric for path selection,
- exchanges RIP broadcasts with neighboring devices to dynamically learn and advertise routes, and
- includes four basic components: routing update process, RIP routing metrics, routing stability, and routing timers.

RIP characteristics and versions

RIP is a widely used routing protocol. It sends routing updates at regular intervals and when the network topology changes. These updates provide information about network reachability and the number of hops to each destination..

You can configure RIP with less effort than OSPF, but it generates more network traffic. RIP supports both Version 1 and Version 2.

Secure Firewall Threat Defense device supports both RIP Version 1 and Version 2. RIP Version 1 does not send subnet masks with its updates; Version 2 supports variable-length subnet masks and neighbor authentication to ensure trusted routing information exchanges.

RIP offers simple configuration and adapts to topology changes. In contrast, static routes require manual updates for each change. RIP, however, incurs higher network and processing overhead compared to static routing.

How RIP routing updates work

Summary

Key components in the RIP routing update process include:

- RIP routers: Send and receive routing update messages to maintain network topology information
- Routing tables: Store the best routes to destinations with associated metric values
- Routing update messages: Contain network topology changes and route information
- Metric values: Determine the best path, with routers maintaining only the lowest metric route

Workflow

The process involves these stages:

1. RIP routers send routing-update messages at regular intervals and when the network topology changes.
2. When a router receives a routing update that includes changes to an entry, it updates its routing table to reflect the new route. The metric value for the path is increased by 1, and the sender is indicated as the next hop.
3. RIP routers maintain only the best route to a destination, which is the route with the lowest metric value.
4. After updating its routing table, the router immediately begins transmitting routing updates to inform other network routers of the change. These updates occur independently from regularly scheduled RIP router updates.

RIP routing metric

A RIP routing metric is a network measurement that

- calculates the distance between a source and a destination network,
- assigns a hop count value, typically 1, to each hop in a network path from the source to the destination, and
- updates the routing table when new routing data is received, using the sender's IP address as the next hop.

RIP stability features

RIP stability features are routing mechanisms that

- prevent routing loops by setting a limit on the number of hops in a path from source to destination,
- provide stability despite potentially rapid changes in network topology, and
- include split horizon and hold-down mechanisms to prevent incorrect routing information propagation.

Additional stability mechanisms

The maximum number of hops allowed in a path is 15. If a routing update shows a new or changed entry, the system increases the metric value to infinity (16), which makes the network destination unreachable. This stability mechanism limits the maximum diameter of a RIP network to fewer than 16 hops.

RIP timers

RIP timers are network protocol mechanisms that

- regulate RIP performance through timed intervals for routing updates and route management,
- control when routes are marked invalid, placed in holddown, or flushed from the routing table, and
- help prevent network congestion by managing the timing of routing updates across devices.

RIP timer stages

RIP uses numerous timers to regulate its performance. These are the timer stages for RIP:

- Update: The routing-update timer is the interval between periodic routing updates. This is how often the device sends routing updates. Generally, set to 30 seconds, a small random delay added to prevent simultaneous route updates.
- Invalid: Each routing table entry has a route-timeout timer for the last valid update. The route is marked invalid and moves to holddown once the timer expires. The default is 180 seconds (3 minutes).
- Holddown: The period the system waits before accepting any new updates for invalid routes. The default is 180 seconds (3 minutes).
- Flush: The route-flush timer marks the period from when the system receives the last valid update until the route is discarded and removed from the routing table. The default is 240 seconds (4 minutes).

RIP timer sequence when adjacent router interface goes down

When an adjacent router goes down, routing updates stop, and the Invalid and Flush timers start increasing. In the first 180 seconds, nothing will happen. After 180 seconds, the Invalid timer expires, marks the route invalid, and the Holddown timer begins, holding the route for 60 seconds.

If there is no routing update for the adjacent router's status (that is, it is still down), the route enters into the Flush state. In total, the system has waited for 240 seconds from the last update (180 seconds for the Invalid timer and 60 seconds for Holddown timer), and then the system flushes the route. Even if the adjacent router's interface comes up immediately, the system waits for the Holddown timer to complete the 120-second period before accepting updates.

Requirements and prerequisites for RIP

These requirements and prerequisites must be met for configuring RIP on supported firewall models.

Model support

Firewall Threat Defense

Firewall Threat Defense Virtual

Supported domains

Any

User roles

Admin

Network Admin

Guidelines for RIP

IPv6 guidelines

RIP does not support IPv6.

RIP version 2 guidelines

Consider these guidelines for RIP Version 2 configuration:

- If using neighbor authentication, the authentication key and key ID must be the same on all neighbor devices that provide RIP Version 2 updates to the interface.
- With RIP Version 2, the Secure Firewall Threat Defense device transmits and receives default route updates using the multicast address 224.0.0.9. In passive mode, it receives route updates at that address.
- When RIP Version 2 is configured on an interface, the multicast address 224.0.0.9 is registered on that interface. When a RIP Version 2 configuration is removed from an interface, that multicast address is unregistered.

RIP limitations

These are the limitations of RIP:

- The Secure Firewall Threat Defense device cannot pass RIP updates between interfaces.
- RIP Version 1 does not support variable-length subnet masks.
- RIP has a maximum hop count of 15. A route with a hop count greater than 15 is considered unreachable.
- RIP convergence is relatively slow compared to other routing protocols.
- You can only enable a single RIP process on the Secure Firewall Threat Defense device.

Configure RIP

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection.

Use this procedure to define networks, manage redistribution, and apply interface-specific filtering or broadcast parameters.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** and select **RIP** from the table of contents.

Check the **Enable RIP** check box to configure the RIP settings.

Choose the RIP versions for sending and receiving RIP updates from the **RIP Version** drop-down list:

- **Receive Version 1 & 2, Send Version 1**
- **Send & Receive Version 1**
- **Send & Receive Version 2**

Step 3 (Optional) Check the **Generate Default Route** check box to generate a default route for distribution, based on the route map that you specify.

a) Specify a route map name to use for generating default routes, in the **Route Map** field.

Generate the default route 0.0.0.0/0 for distribution when a route map is specified in the **Route Map** field.

Step 4 When you choose Send and Receive Version 2, the **Enable Auto Summary** option is available. When the **Enable Auto Summary** check box is checked, automatic route summarization is enabled. Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is disabled, subnets are advertised.

Note
RIP Version 1 always uses automatic summarization—you cannot disable it.

Step 5 Click **Networks**. Define one or more networks for RIP routing. Enter IP address(es), or enter or select the desired Network/Hosts objects. There is no limit to the number of networks you can add to the security appliance configuration. Interfaces belonging to the defined network will participate in the RIP routing. The RIP routing updates will be sent and received only through interfaces on the specified networks. Also, if the network of an interface is not specified, the interface will not be advertised in any RIP updates.

Note
RIP only supports IPv4 objects.

Step 6 (Optional) Click **Passive Interface**. Use this option to specify passive interfaces on the appliance, and by extension the active interfaces. The device receives RIP routing broadcasts on passive interfaces, using that information to populate its routing tables, but does not broadcast routing updates on passive interfaces. Interfaces that are not designated as passive receive and send updates.

Step 7 Click **Redistribution** to manage redistribution routes. These are the routes that are being redistributed from other routing processes into the RIP routing process.

a) Click **Add** to specify redistribution routes.
b) Choose the routing protocol to redistribute into the RIP routing process, in the **Protocol** drop-down list. Available options are **Static**, **Connected**, **OSPF**, **BGP**, and **EIGRP**.

Note
For OSPF, specify the process ID in the **Process ID** field. For BGP, specify the AS number in the **AS Number** field. When you choose the Connected option in the **Protocol** drop-down list, you can redistribute directly connected networks into the RIP routing process.

c) (Optional) If you are redistributing OSPF routes into the RIP routing process, you can select specific types of OSPF routes to redistribute in the **Match** drop-down list. Ctrl-click to select multiple types:

- **Internal**: Routes internal to the autonomous system (AS) are redistributed.
- **External 1**: Type 1 routes external to the AS are redistributed.
- **External 2**: Type 2 routes external to the AS are redistributed.
- **NSSA External 1**: Type 1 routes external to a not-so-stubby area (NSSA) are redistributed.
- **NSSA External 2**: Type 2 routes external to an NSSA are redistributed.

Note
The default is match Internal, External 1, and External 2.

d) Select the RIP metric type to apply to the redistributed routes in the **Metric** drop-down list. The three choices are:

- **Transparent**: Use the current route metric.
- **Specified Value**: Assign a specific metric value from 0 to 16 in the **Metric Value** field.
- **None**: No metric is specified. Do not use any metric value, to apply to redistributed routes.

Note
None option is applicable only for **Static** and **Connected** protocols.

e) (Optional) Enter the name of a route map that must be satisfied, in the **Route Map** field before the route can be redistributed into the RIP routing process. Routes are redistributed only if IP address matches an allow statement in the route map address list. To create a new route map object, click **Add** (+). See Configure Route Map Entry for the procedure to add a new route map.
f) Click **OK**.

Step 8 (Optional) Click **Filtering** to manage filters for the RIP policy. In this section, filters are used to prevent routing updates through an interface, control the advertising of routes in routing updates, control the processing of routing updates and filtering sources of routing updates.

a) Click **Add** to add RIP filters.
b) Select the type of traffic to be filtered - Inbound or Outbound in the **Traffic Direction** field.

Note
If traffic direction is inbound, you can only define an Interface filter.

c) Specify whether the filter is based on an Interface or a Route, by selecting appropriate option in the **Filter On** field. If you click **Interface**, you can enter or choose the name of the interface on which routing updates are to be filtered. If you click **Route**, you can choose the route type:

- **Static**: Only static routes are filtered.
- **Connected**: Only connected routes are filtered.
- **OSPF**: Only OSPFv2 routes discovered by the specified OSPF process are filtered. Enter the Process ID of the OSPF process to be filtered.
- **BGP**: Only BGPv4 routes discovered by the specified BGP process are filtered. Enter the value in the **Process ID** field.
- **EIGRP**: Only EIGRP routes are filtered.

d) In the **Access List** field, enter or choose the name of one or more access control lists (ACLs) that define the networks to be allowed or removed from RIP route advertisements. To add a new standard access list object, click **Add** (+) and see Configure a standard ACL object, on page 2953.

e) Click **OK**.

Step 9 (Optional) Click **Broadcast** to add or edit interface configurations. Using Broadcast, you can override the global RIP versions to send or receive per interface. You can also define the authentication parameters per interface if you want to implement authentication to ensure valid RIP updates.

a) Click **Add** to add interface configurations.

b) Enter or choose an interface defined on this appliance in the **Interface** field.

c) In the Send option, select the appropriate boxes to specify sending updates using the RIP **Version 1**, **Version 2**, or both. These options let you override, for the specified interface, the global Send versions specified.

d) In the Receive option, select the appropriate boxes to specify accepting updates using the RIP **Version 1**, **Version 2**, or both. These options let you override, for the specified interface, the global Receive versions specified.

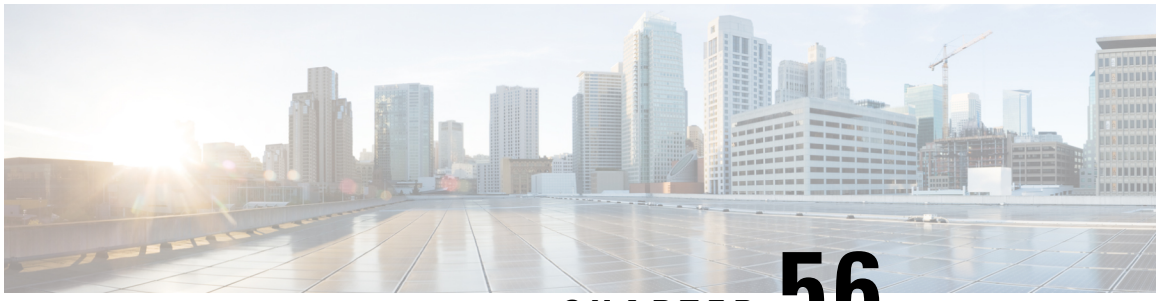
e) Select the **Authentication** used on this interface for RIP broadcasts.

- **None**: No authentication
- **MD5**: Employ MD5
- **Clear Text**: Employ clear-text authentication

If you choose MD5 or Clear Text, you must also provide the following authentication parameters.

- **Key ID**: The ID of the authentication key; valid values are from 0 to 255.
- **Key**: The key used by the chosen authentication method; can contain up to 16 characters.
- **Confirm**: Enter the authentication key again to confirm.

f) Click **OK**.



CHAPTER 56

Multicast

This chapter describes how to configure the Secure Firewall Threat Defense device to use the multicast routing protocol.

- Multicast routing, on page 1521
- Prerequisites for Multicast Routing, on page 1526
- Guidelines for multicast routing, on page 1526
- Configure IGMP features, on page 1527
- Configure PIM features, on page 1532
- Configure multicast routes, on page 1538
- Configure multicast boundary filters, on page 1539

Multicast routing

Multicast routing is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients and homes. Applications include videoconferencing, corporate communications, distance learning, and distributing software, stock quotes, and news.

Multicast routing applications and protocols

Multicast routing protocols send source traffic to multiple receivers efficiently, without adding extra load on the source or receivers, and use less network bandwidth than other technologies. Multicast packets are replicated in the network by Firewall Threat Defense device enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols, which results in the most efficient delivery of data to multiple receivers possible.

The Firewall Threat Defense device supports both stub multicast routing and PIM multicast routing. However, you cannot configure both concurrently on a single Firewall Threat Defense device.

Note The UDP and non-UDP transports are both supported for multicast routing. However, the non-UDP transport has no FastPath optimization.

IGMP Protocol

IP hosts use the Internet Group Management Protocol (IGMP) to report their group memberships to directly-connected multicast routers. IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

IGMP uses group addresses (Class D IP address) as group identifiers. Host group address can be in the range of 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is never assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.

Note When you enable multicast routing on the Firewall Threat Defense device, IGMP Version 2 is automatically enabled on all interfaces.

Query Messages to Multicast Groups

The Firewall Threat Defense device sends query messages to discover which multicast groups have members on the networks attached to the interfaces. Members respond with IGMP report messages indicating that they want to receive multicast packets for specific groups. Query messages are addressed to the all-systems multicast group, which has an address of 224.0.0.1, with a time-to-live value of 1.

These messages are sent periodically to refresh the membership information stored on the Firewall Threat Defense device. If the Firewall Threat Defense device discovers that there are no local members of a multicast group still attached to an interface, it stops forwarding multicast packets for that group to the attached network, and it sends a prune message back to the source of the packets.

By default, the PIM designated router on the subnet is responsible for sending the query messages. By default, they are sent once every 125 seconds.

When changing the query response time, by default, the maximum query response time advertised in IGMP queries is 10 seconds. If the Firewall Threat Defense device does not receive a response to a host query within this amount of time, it deletes the group.

Stub multicast routing

Stub multicast routing provides dynamic host registration and facilitates multicast routing.

When configured for stub multicast routing, the Firewall Threat Defense device acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the Firewall Threat Defense device forwards IGMP messages to an upstream multicast router, which sets up delivery of the multicast data. When configured for stub multicast routing, the Firewall Threat Defense device cannot be configured for PIM sparse or bidirectional mode. Enable PIM on the interfaces that participate in IGMP stub multicast routing.

The Firewall Threat Defense device supports both PIM-SM and bidirectional PIM. PIM-SM is a multicast routing protocol that uses the underlying unicast routing information base or a separate multicast-capable routing information base. It builds unidirectional shared trees, each rooted at a single Rendezvous Point (RP) per multicast group. Optionally, it creates shortest-path trees for each multicast source.

PIM multicast routing

Bidirectional PIM is a variant of PIM-SM that builds bidirectional shared trees connecting multicast sources and receivers.

DF election process

Bidirectional trees are built using a Designated Forwarder (DF) election process operating on each link of the multicast topology. With the assistance of the DF, multicast data is forwarded from sources to the Rendezvous Point (RP) and then along the shared tree to receivers. This process does not require source-specific state. The DF election takes place during RP discovery and provides a default route to the RP.



Note If the Firewall Threat Defense device is the PIM RP, use the untranslated outside address of the Firewall Threat Defense device as the RP address.

PIM sparse mode

PIM Sparse Mode (PIM-SM) is a PIM mode that forwards traffic only when specifically requested by receivers. It uses a central Rendezvous Point (RP) to coordinate the multicast traffic forwarding between the multicast source and the destination. RP is a router that serves as the meeting point for the source and the receivers. Multicast sources register with the RP through their first hop router. They send a copy of the multicast traffic to the receivers through the RP.

PIM source specific multicast support

Firewall Threat Defense device can forward source specific multicast traffic (e.g., for groups in the 232.x.x.x range) even though they do not support PIM-SSM configuration or Internet Group Management Protocol version 3 (IGMPv3).

SSM is classified as a data delivery mechanism for one-to-many applications such as IPTV. The SSM model uses a concept of "channels" denoted by an (S,G) pair, where S is a source address and G is an SSM destination address. Subscribing to a channel is achieved by using a group management protocol such as IGMPv3. SSM enables a receiving client, once it has learned about a particular multicast source, to receive multicast streams directly from the source rather than receiving it from a shared Rendezvous Point (RP). Access control mechanisms are introduced within SSM, providing a security enhancement that is not available with current sparse or sparse-dense mode implementations.

Limitation:

- Firewall Threat Defense device cannot act as the last-hop router for SSM multicast because it does not support IGMPv3 (which receivers use to join SSM groups).

  If the Firewall Threat Defense device is the last-hop, it will ignore IGMPv3 join messages from receivers for SSM groups, and SSM forwarding will not work.

- Static multicast routes on Firewall Threat Defense device do not work for SSM range (232.x.x.x).

Workaround:

For Cisco Firewall Threat Defense device firewall to forward SSM-related multicast traffic, you must add a multicast-capable layer 3 device (such as a router or switch) that supports PIM and IGMPv3 on the same network segment as the receivers.

How it Works:

- Receivers register their SSM group joins (IGMPv3) with this layer 3 device.
- This layer 3 device sends PIM join messages towards the Firewall Threat Defense device.
- The Firewall Threat Defense device receives these PIM messages and dynamically learns multicast routes.
- SSM multicast traffic is then properly forwarded by the Firewall Threat Defense device, since it is no longer the last-hop device.

Multicast Bidirectional PIM

Multicast bidirectional PIM is useful for networks that have many sources and receivers talking to each other simultaneously and where each participant can become both the source and receiver of multicast traffic, such as in videoconferencing, Webex meetings, and group chat. When PIM bidirectional mode is used, the RP only creates the (*,G) entry for the shared tree. There is no (S,G) entry. This conserves resources on the RP because state tables for each (S,G) entry are not maintained.

In PIM sparse mode, traffic only flows down the shared tree. In PIM bidirectional mode, traffic flows up and down the shared tree.

PIM bidirectional mode also does not use the PIM register/register-stop mechanism to register sources to the RP. Each source can begin sending to the source at any time. When the multicast packets arrive at the RP, they are forwarded down the shared tree (if there are receivers) or dropped (when there are no receivers). However, there is no way for the RP to tell the source to stop sending multicast traffic.

Design-wise you must think about where to place the RP in your network because it should be somewhere in the middle between the sources and receivers in the network.

PIM bidirectional mode has no Reverse Path Forwarding (RPF) check. Instead it uses the concept of a Designated Forwarder (DF) to prevent loops. This DF is the only router on the segment that is allowed to send multicast traffic to the RP. If there is only one router per segment that forwards multicast traffic, there will be no loops. The DF is chosen using the following mechanism:

- The router with the lowest metric to the RP is the DF.
- If the metric is equal, then the router with the highest IP address becomes the DF.

PIM bootstrap router (BSR)

PIM Bootstrap Router (BSR) is a dynamic Rendezvous Point (RP) selection model that uses candidate routers for RP function and for relaying the RP information for a group.

The RP function includes RP discovery and provides a default route to the RP. It does this by configuring a set of devices as candidate BSRs (C-BSR) which participate in a BSR election process to choose a BSR amongst themselves. Once the BSR is chosen, devices that are configured as candidate Rendezvous Points (C-RP) start sending their group mapping to the elected BSR. The BSR then distributes the group-to-RP mapping information to all the other devices down the multicast tree through BSR messages that travel from PIM router to PIM router on a per-hop basis.

This feature allows networks to learn RPs dynamically. Dynamic RP discovery is essential in large, complex networks where an RP may occasionally go down and come up.

PIM bootstrap router (BSR) terminology

This reference provides definitions for key terminology used in PIM Bootstrap Router (BSR) configurations to help understand the roles and functions of different components in the BSR election and RP advertisement process.

These terms are frequently referenced in the PIM BSR configuration:

- Bootstrap Router (BSR) — A BSR advertises Rendezvous Point (RP) information to other routers with PIM on a hop-by-hop basis. Among multiple Candidate-BSRs, a single BSR is chosen after an election process. The primary purpose of this Bootstrap router is to collect all Candidate-RP (C-RP) announcements in to a database called the RP-set and to periodically send this out to all other routers in the network as BSR messages (every 60 seconds).
- Bootstrap Router (BSR) messages — BSR messages are multicast to the All-PIM-Routers group with a TTL of 1. All PIM neighbors that receive these messages retransmit them (again with a TTL of 1) out of all interfaces except the one in which the messages were received. BSR messages contain the RP-set and the IP address of the currently active BSR. This is how C-RPs know where to unicast their C-RP messages.
- Candidate Bootstrap Router (C-BSR) — A device that is configured as a candidate-BSR participates in the BSR election mechanism. A C-BSR with highest priority is elected as the BSR. The highest IP address of the C-BSR is used as a tiebreaker. The BSR election process is preemptive, for example if a new C-BSR with a higher priority comes up, it triggers a new election process.
- Candidate Rendezvous Point (C-RP) — An RP acts as a meeting place for sources and receivers of multicast data. A device that is configured as a C-RP periodically advertises the multicast group mapping information directly to the elected BSR through unicast. These messages contain the Group-range, C-RP address, and a hold time. The IP address of the current BSR is learned from the periodic BSR messages that are received by all routers in the network. In this way, the BSR learns about possible RPs that are currently up and reachable.



Note The Firewall Threat Defense device does not act as a C-RP, even though the C-RP is a mandatory requirement for BSR traffic. Only routers can act as a C-RP. So, for BSR testing functionality, you must add routers to the topology.

- BSR Election Mechanism — Each C-BSR originates Bootstrap messages (BSMs) that contain a BSR Priority field. Routers within the domain flood the BSMs throughout the domain. A C-BSR that hears about a higher-priority C-BSR than itself suppresses its sending of further BSMs for some period of time. The single remaining C-BSR becomes the elected BSR, and its BSMs inform all the other routers in the domain that it is the elected BSR.

Multicast group concept

Multicast is based on the concept of a group. An arbitrary group of receivers expresses an interest in receiving a particular data stream. This group does not have any physical or geographical boundaries. Hosts can be

located anywhere on the Internet. Hosts that want to receive data sent to a group must join the group using IGMP.

Group membership requirements

Hosts must be a member of the group to receive the data stream.

Multicast Addresses

Multicast addresses specify an arbitrary group of IP hosts that have joined the group and want to receive traffic sent to this group.

Clustering

Multicast routing supports clustering. In Spanned EtherChannel clustering, the control unit sends all multicast routing packets and data packets until fast-path forwarding is established.

After fast-path forwarding is established, data units may forward multicast data packets. All data flows are full flows. Stub forwarding flows are also supported. Because only one unit receives multicast packets in Spanned EtherChannel clustering, redirection to the control unit is common.

Prerequisites for Multicast Routing

Model support

Firewall Threat Defense

Firewall Threat Defense Virtual

Supported domains

Any

User roles

- Admin
- Network Admin

Guidelines for multicast routing

Table 95: System requirements

Item	Description
Context mode	Supported in single context mode.
Firewall mode	Supported only in routed firewall mode. Transparent firewall mode is not supported.

Item	Description
IPv6	IPv6 is not supported.
Multicast group	The range of addresses from 224.0.0.0 to 224.0.0.255 is reserved for routing protocols and other topology discovery or maintenance protocols, such as gateway discovery and group membership reporting. Internet multicast routing from the 224.0.0/24 address range is not supported. When multicast routing is enabled for reserved addresses, an IGMP group is not created.
Clustering	In clustering, for IGMP and PIM, this feature is only supported on the primary unit.

Access control configuration

You must configure an access control or prefilter rule on the inbound security zone to allow traffic to the multicast host, such as 224.1.2.3. However, you cannot specify a destination security zone for the rule, or it cannot be applied to multicast connections during initial connection validation.

Interface and protocol management

Ensure proper configuration management when working with PIM-enabled interfaces and multicast routing protocols.

- You cannot disable an interface with PIM configured on it. If you have configured PIM on the interface (see Configure PIM protocol, on page 1533), disabling the multicast routing and PIM does not remove the PIM configuration. You must remove (delete) the PIM configuration to disable the interface.
- Multicast routing is not supported in ECMP.
- Do not configure Firewall Threat Defense to simultaneously be a Rendezvous Point (RP) and a First Hop Router.
- The HSRP standby IP address does not participate in PIM neighborship. If the RP router IP is routed through the HSRP standby IP address, multicast routing does not work in Firewall Threat Defense. To ensure multicast traffic passes through successfully, configure the route for the RP address to use an interface IP address rather than the HSRP standby IP address.
- For a device using virtual routing, you can configure multicast only for its global virtual router and not for its user-defined virtual router.

Configure IGMP features

Configure IGMP features to enable dynamic registration of individual hosts in multicast groups on a particular LAN and to manage multicast traffic efficiently.

IP hosts use IGMP to report their group memberships to directly-connected multicast routers. IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

IP hosts use IGMP to report their group memberships to directly connected multicast routers. IGMP dynamically registers individual hosts in a multicast group on a specific LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

This section describes how to configure optional IGMP settings on a per-interface basis.

Procedure

Step 1 Enable multicast routing, on page 1528.

Step 2 Configure IGMP protocol, on page 1529.

Step 3 Configure IGMP access groups, on page 1530.

Step 4 Configure IGMP static groups, on page 1531.

Step 5 Configure IGMP join groups, on page 1531.

Enable multicast routing

Enable multicast routing to allow the device to forward multicast traffic using Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM) protocols on all interfaces.

Enabling multicast routing on the Firewall Threat Defense device, enables IGMP and PIM on all interfaces by default. IGMP is used to learn whether members of a group are present on directly attached subnets. Hosts join multicast groups by sending IGMP report messages. PIM is used to maintain forwarding tables to forward multicast datagrams.

Note Only the UDP transport layer is supported for multicast routing.

These are the maximum number of entries for specific multicast tables. Any new entries are discarded when these limits are reached.

- Multicast Forwarding Information Base (MFIB): 30,000
- IGMP Groups: 30,000
- PIM Routes: 72,000

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 Check the **Enable Multicast Routing** check box.

Checking this check box enables IP multicast routing on the device. Unchecking it disables IP multicast routing. By default, multicast is disabled. Enabling multicast routing enables multicast on all interfaces.

You can disable multicast on a per-interface basis. This is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the Firewall Threat Defense device from sending host query messages on that interface.

Configure IGMP protocol

Use this task to configures IGMP parameters per-interface. These settings help manage multicast group membership on a Firewall Threat Defense device.

You can configure IGMP parameters per interface, such as the forward interface, query messages, and time intervals.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 Click the **Protocol** tab and click **Add**, or edit an interface.

In the **Add IGMP parameters** dialog box, configure new IGMP parameters to the Firewall Threat Defense device. Use the **Edit IGMP parameters** dialog box to modify the existing parameters.

Step 4 Configure these parameters:

- **Interface**: From the drop-down list, choose the interface for which you want to configure IGMP protocol.
- **Enable IGMP**: Check the check box to enable IGMP.

  Note
  Disabling IGMP on specific interfaces is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the device from sending host query messages on that interface.

- **Forward Interface**: From the drop-down list, choose the specific interface from which you want to forward IGMP messages.

  This configures the Secure Firewall Threat Defense device to act as an IGMP proxy agent and forward IGMP messages from hosts connected on one interface to an upstream multicast router on another interface.

- **Version**: Choose IGMP Version 1 or 2.

  By default, the Firewall Threat Defense device runs IGMP Version 2, which enables several additional features.

  Note
  All multicast routers on a subnet must support the same version of IGMP. The Firewall Threat Defense device does not automatically detect Version 1 routers and switch to Version 1. However, you can have a mix of IGMP Version 1 and 2 hosts on the subnet; the Firewall Threat Defense device running IGMP Version 2 works correctly when IGMP Version 1 hosts are present.

- **Query Interval**: The interval in seconds at which the designated router sends IGMP host-query messages. The range is 1 to 3600. The default is 125.

Note
If the Firewall Threat Defense device does not hear a query message on an interface for the specified timeout value, then the device becomes the designated router and starts sending the query messages.

- **Response Time**: The interval in seconds before the Firewall Threat Defense device deletes the group. The range is 1 to 25. The default is 10.

  If the Firewall Threat Defense device does not receive a response to a host query within this amount of time, it deletes the group.

- **Group Limit**: Maximum number of hosts that can join on an interface. The range is 1 to 500. The default is 500.

  You can limit the number of IGMP states resulting from IGMP membership reports on a per-interface basis. Membership reports that exceed the configured limits are not entered in the IGMP cache, and their traffic is not forwarded.

- **Query Timeout**: The period of time in seconds before which the Firewall Threat Defense device takes over as the requester for the interface after the previous requester has stopped. The range is 60 to 300. The default is 255.

Step 5 Click **OK** to save the IGMP protocol configuration.

Configure IGMP access groups

You can control access to multicast groups by using access control lists.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Access Group**.

Step 3 On **Access Group**, click **Add** or **Edit**.

Use the **Add IGMP Access Group parameters** dialog box to add new IGMP access groups to the Access Group table. Use the **Edit IGMP Access Group parameters** dialog box to change the parameters of an existing access group.

Step 4 Configure the access group options:

a) From the **Interface** drop-down list, choose the interface associated with the access group. You cannot change the associated interface when you are editing an existing access group.

b) Click one of the following:

- **Standard Access List**: From the **Standard Access List** drop-down list, choose the standard ACL or click **Add** (+) to create a new standard ACL. For more information, refer to Configure a standard ACL object, on page 2953.
- **Extended Access List**: From the **Extended Access List** drop-down list, choose the extended ACL or click **Add** (+) to create a new extended ACL. For more information, refer to Configure an extended ACL object, on page 2949.

Step 5 Click **OK** to save the access group configuration.

Configure IGMP static groups

Configure a statically joined IGMP group to enable multicast traffic for that group to be sent to a network segment even when group members cannot report membership or no members exist on the segment.

Sometimes a group member cannot report its membership in the group or there may be no members of a group on the network segment, but you still want multicast traffic for that group to be sent to that network segment. You can have multicast traffic for that group sent to the segment by configuring a statically joined IGMP group. With this method, the Firewall Threat Defense device does not accept the packets itself, but only forwards them. Therefore, this method allows fast switching. The outgoing interface appears in the IGMP cache, but this interface is not a member of the multicast group.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Static Group**, click **Add** or **Edit**.

Use the **Add IGMP Static Group parameters** dialog box to statically assign a multicast group to an interface. Use the **Edit IGMP Static Group parameters** dialog box to change existing static group assignments.

Note
The IGMP Static Group enables PIM to send *Join* requests towards the sources or towards the Rendezvous Point (RP), provided, the firewall with this command is the PIM Designated Router (DR) on that interface where the command is applied.

Step 4 Configure the following options:

- From the **Interface** drop-down list, choose the interface to which you want to statically assign a multicast group. If you are editing an existing entry, you cannot change the value.
- From the **Multicast Groups** drop-down list, choose the multicast group to which you want to assign the interface, or click **Add** (+) to create a new multicast group. Refer to Creating Network Objects for the procedure.

Step 5 Click **OK** to save the static group configuration.

Configure IGMP join groups

You can configure an interface to be a member of a multicast group.

Configuring the Firewall Threat Defense device to join a multicast group, upstream routers maintain multicast routing table information for that group and keep the group's paths active.

Note Refer to Configure IGMP static groups, on page 1531 if you want to forward multicast packets for a specific group to an interface without the Firewall Threat Defense device accepting those packets as part of the group.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Join Group**, click **Add** or **Edit**.

Use the **Add IGMP Join Group parameters** dialog box to configure the Firewall Threat Defense device to be a member of a multicast group. Use the **Edit IGMP Join Group parameters** dialog box to change existing parameters.

Note
The IGMP Join Group enables PIM to send *Join* requests to sources or to the Rendezvous Point (RP) if the firewall with this command is the PIM Designated Router (DR) on that interface.

Step 4 Configure the following options:

- From the **Interface** drop-down list, choose the interface you want to be a member of a multicast group. If you are editing an existing entry, you cannot change the value.
- From the **Join Group** drop-down list, choose the multicast group to which you want to assign the interface, or click **Plus** to create a new multicast group. Refer to Creating Network Objects for the procedure.

Configure PIM features

This section describes how to configure optional PIM settings.

Routers use PIM to maintain forwarding tables to use for forwarding multicast diagrams. Enable multicast routing on the Secure Firewall Threat Defense device, to automatically enable IGMP and PIM on all interfaces.

Note PIM is not supported with PAT. The PIM protocol does not use ports, and PAT only works with protocols that use ports.

Procedure

Step 1 Configure PIM protocol, on page 1533.

Step 2 Configure PIM neighbor filters, on page 1534.

Step 3 Configure PIM bidirectional neighbor filters, on page 1534.

Step 4 Configure PIM rendezvous points, on page 1535.

Step 5 Configure PIM route trees, on page 1536.

Step 6 Configure PIM request filters, on page 1537.

Step 7 Configure multicast boundary filters, on page 1539.

Configure PIM protocol

You can enable or disable PIM on a specific interface.

You can also configure the Designated Router (DR) priority. The DR is responsible for sending PIM register, join, and prune messages to the RP. When there is more than one multicast router on a network segment, the DR is chosen based on DR priority. If multiple devices share the same DR priority, the device with the highest IP address becomes the DR. By default, the Firewall Threat Defense device has a DR priority of 1.

Router query messages are used to choose the PIM DR. The PIM DR is responsible for sending router query messages. By default, router query messages are sent every 30 seconds. Additionally, every 60 seconds, the Firewall Threat Defense device sends PIM join or prune messages.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 Click the **Protocol** tab and click **Add**, or edit an interface.

In the **Add PIM parameters** dialog box, add new PIM parameters for the interface. Use the **Edit PIM parameters** dialog box to change existing parameters.

Step 4 Configure these options:

- **Interface**: From the drop-down list, select the interface for which you want to configure PIM protocol.
- **Enable PIM**: Check the check box to enable PIM.
- **DR Priority**: Set the DR priority value for the selected interface. The router with the highest DR priority on the subnet becomes the designated router.
- **Hello Interval**: The interval in seconds at which the interface sends PIM hello messages. The range is 1 to 3600. The default is 30.
- **Join Prune Interval**: The interval in seconds at which the interface sends PIM join and prune advertisements. The range is 10 to 600. The default is 60.

Step 5 Click **OK** to save the PIM protocol configuration.

Configure PIM neighbor filters

You can define and filter the routers that become Protocol Independent Multicast (PIM) neighbors.

PIM neighbor filters allow you to define and filter the routers that become Protocol Independent Multicast (PIM) neighbors to:

- Prevent unauthorized routers from becoming PIM neighbors.
- Prevent attached stub routers from participating in PIM.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 Click the **Neighbor Filter** tab and click **Add**, or edit an interface.

In the **Add PIM Neighbor Filter** dialog box, add new PIM neighbor filters to the interface. Use the **Edit PIM Neighbor Filter** dialog box to change existing parameters.

Step 4 Configure these options:

- From the **Interface** drop-down list, choose the interface to which you want to add a PIM neighbor filter.
- From the **Standard Access List** drop-down list, choose a standard ACL or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object, on page 2953 for the procedure.

  Note
  Choosing **Allow** on the **Add Standard Access List Entry** dialog box lets the multicast group advertisements pass through the interface. Choosing **Block** prevents the specified multicast group advertisements from passing through the interface. When a multicast boundary is configured on an interface, all multicast traffic is prevented from passing through the interface unless permitted with a neighbor filter entry.

Step 5 Click **OK** to save the PIM neighbor filter configuration.

Configure PIM bidirectional neighbor filters

Configure PIM bidirectional neighbor filters to control which neighbor devices participate in the Designated Forwarder (DF) election for bidirectional multicast routing. This helps ensure proper multicast traffic forwarding in the network.

A PIM bidirectional neighbor filter is an ACL that defines the neighbor devices that can participate in the Designated Forwarder (DF) election. If you do not configure a PIM bidirectional neighbor filter for an interface, there are no restrictions. If you configure a PIM bidirectional neighbor filter, only neighbors permitted by the ACL can participate in the DF election process.

Bidirectional PIM allows multicast routers to keep reduced state information. All multicast routers in a segment must be bidirectionally enabled to elect a DF.

When a PIM bidirectional neighbor filter is enabled, the routers permitted by the ACL are considered to be bidirectionally capable. Therefore, the following is true:

- If a permitted neighbor does not support bidirectional mode, the DF election does not occur.
- If a denied neighbor supports bidirectional mode, the DF election does not occur.
- If a denied neighbor does not support bidirectional mode, the DF election can occur.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Multicast Routing** > **PIM**.

Step 3 On **Bidirectional Neighbor Filter**, click **Add** or **Edit**.

Use the **Add PIM Bidirectional Neighbor Filter** dialog box to create ACL entries for the PIM bidirectional neighbor filter ACL. Use the **Edit PIM Bidirectional Neighbor Filter** dialog box to change existing parameters.

Step 4 Configure these options:

- In the **Interface** drop-down list, select the interface for which you want to configure the PIM bidirectional neighbor filter ACL entry.
- From the **Standard Access List** drop-down list, select a standard ACL or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object, on page 2953 for the procedure.

  Note
  Choosing **Allow** in the **Add Standard Access List Entry** dialog box lets the specified devices participate in the DR election process. Choosing **Block** prevents the specified devices from participating in the DR election process.

Step 5 Click **OK** to save the PIM bidirectional neighbor filter configuration.

Configure PIM rendezvous points

You can configure the Firewall Threat Defense device to serve as a RP to more than one group.

The group range specified in the ACL determines the PIM RP group mapping. If an ACL is not specified, then the RP for the group is applied to the entire multicast group range (224.0.0.0/4). Refer to Multicast Bidirectional PIM, on page 1524 for more information about bidirectional PIM.

These restrictions apply to RPs:

- You cannot use the same RP address twice.
- You cannot specify **All Groups** for more than one RP.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Rendezvous Points**, click **Add** or **Edit**.

Use the **Add Rendezvous Point** dialog box to create a new entry in the Rendezvous Point table. Use the **Edit Rendezvous Point** dialog box to modify existing parameters.

Step 4 Configure these options:

- From the **Rendezvous Point IP address** drop-down list, choose the IP address that you want to add as an RP or click **Add** (+) to create a new network object. Refer to Creating Network Objects for the procedure.
- Check the **Use bi-directional forwarding** check box to enable bidirectional mode for the specified multicast groups. In this mode, if the Firewall Threat Defense device receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a prune message back to the source.
- Click **Use this RP for all Multicast Groups** to use the specified RP for all multicast groups on the interface.
- Click **Use this RP for all Multicast Groups as specified below** to designate the multicast groups to use with the specified RP and then from the **Standard Access List** drop-down list, choose a standard ACL or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object, on page 2953 for the procedure.

Step 5 Click **OK** to save the rendezvous point configuration.

Configure PIM route trees

Configure PIM route trees to control whether the Firewall Threat Defense device joins the shortest-path tree or uses the shared tree for multicast traffic routing.

By default, PIM leaf routers join the shortest-path tree immediately after the first packet arrives from a new source. This method reduces delay, but requires more memory than the shared tree. You can configure whether or not the Firewall Threat Defense device should join the shortest-path tree or use the shared tree, either for all multicast groups or only for specific multicast addresses.

The shortest-path tree is used for any group that is not specified in the Multicast Groups table. The Multicast Groups table displays the multicast groups to use with the shared tree. The table entries are processed from the top down. You can create an entry that includes a range of multicast groups, but excludes specific groups within that range by placing deny rules for the specific groups at the top of the table and the permit rule for the range of multicast groups below the deny statements.

Note This behavior is known as Shortest Path Switchover (SPT). We recommend that you always use the Shared Tree option.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Route Tree**, select the path for the route tree:

- Click **Shortest Path** to use the shortest-path tree for all multicast groups.
- Click **Shared Tree** to use the shared tree for all multicast groups.
- Click **Shared tree for below mentioned group** to designate the groups specified in the Multicast Groups table, and then from the **Standard Access List** drop-down list, select a standard ACL or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object, on page 2953 for the procedure.

Step 4 Click **OK** to save the route tree configuration.

Configure PIM request filters

When the Firewall Threat Defense device is acting as an RP, you can restrict specific multicast sources from registering with it to prevent unauthorized sources from registering with the RP. You can define the multicast sources from which the Firewall Threat Defense device will accept PIM register messages.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Request Filter**, define the multicast sources that are allowed to register with the Firewall Threat Defense device when it acts as an RP:

- From the **Filter PIM register messages using:** drop-down list select **None**, **Access List**, or **Route Map**.
- If you choose **Access List** from the drop-down list, select an extended ACL or click **Add** (+) to create a new extended ACL. Refer to Configure an extended ACL object, on page 2949 for the procedure.

  Note
  In the **Add Extended Access List Entry** dialog box, select **Allow** from the drop-down list to create a rule that allows the specified source of the specified multicast traffic to register with the Firewall Threat Defense device, or select **Block** to create a rule that prevents the specified source of the specified multicast traffic from registering with the device.

- If you choose **Route Map**, select a route map from the **Route Map** drop-down list, or click **Add** (+) to create a new route map. Refer to Creating Network Objects for the procedure.

Step 4 Click **OK** to save the request filter configuration.

Configure the Firewall Threat Defense device as a candidate bootstrap router

You can configure the Firewall Threat Defense device as a candidate BSR.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Bootstrap Router**, check the **Configure this FTD as a Candidate Bootstrap Router (C-BSR)** check box to perform the C-BSR setup.

a) From the **Interface** drop-down list, select the interface in the Firewall Threat Defense device from which the BSR address is derived to make it a candidate.

This interface must be enabled with PIM.

b) In the **Hash mask length** field, enter the length of a mask (32 bits maximum) that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash (correspond) to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This allows you to configure one RP for multiple groups. The range is 0 to 32.

c) In the **Priority** field, enter the priority of the candidate BSR. The BSR with the larger priority is preferred. If the priority values are the same, the router with the larger IP address is the BSR. The range is 0 to 255. The default value is 0.

Step 4 (Optional) Click **Add** (+) to select an interface on which no PIM BSR messages will be sent or received in the **Configure this FTD as a Border Bootstrap Router (BSR)** section.

- From the **Interface** drop-down list, select the interface on which no PIM BSR messages will be sent or received.

  RP or BSR advertisements are filtered effectively isolating two domains of RP information exchange.

- Check the **Enable Border BSR** check box to enable BSR.

Step 5 Click **OK** to save the bootstrap router configuration.

Configure multicast routes

Configure static multicast routes to separate multicast traffic from unicast traffic. For example, when a path between a source and destination does not support multicast routing, the solution is to configure two multicast devices with a GRE tunnel between them and to send the multicast packets over the tunnel.

When using PIM, the Firewall Threat Defense device expects to receive packets on the same interface where it sends unicast packets back to the source. In some cases, such as bypassing a route that does not support multicast routing, you may want unicast packets to take one path and multicast packets to take another.

Static multicast routes are not advertised or redistributed.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Multicast Routes** > **Add or Edit**.

Use the **Add Multicast Route Configuration** dialog box to add a new multicast route to the Firewall Threat Defense device. Use the **Edit Multicast Route Configuration** dialog box to change an existing multicast route.

Step 3 Choose an existing network from the **Source Network** drop-down, or click **Add** (+) to add a new network.

Refer to Creating Network Objects for the procedure.

Step 4 To configure an interface to forward the route, click **Interface** and configure the following options:

- From the **Source Interface** drop-down list, choose the incoming interface for the multicast route.
- From the **Output Interface/Dense** drop-down list, choose the destination interface that the route is forwarded through.
- In the **Distance** field, enter the distance of the multicast route. The range is 0 to 255.

Step 5 To configure an RPF address to forward the route, click **Address** and configure the following options:

- In the **RPF Address** field, enter the IP address for the multicast route.
- Enter the multicast route distance in the **Distance** field. The value must be between 0 and 255.

Step 6 Click **OK** to save the multicast routes configuration.

Configure multicast boundary filters

Address scoping defines domain boundary filters so that domains with RPs that have the same IP address do not leak into each other. Scoping is performed on the subnet boundaries within large domains and on the boundaries between the domain and the Internet.

You can set up an administratively scoped boundary filter on an interface for multicast group addresses. IANA has designated the multicast address range from 239.0.0.0 to 239.255.255.255 as the administratively scoped addresses. This range of addresses can be reused in domains administered by different organizations. The addresses are considered local and are not globally unique.

A standard ACL defines the range of affected addresses. When a boundary filter is set up, no multicast data packets are allowed to flow across the boundary from either direction. The boundary filter allows the same multicast group address to be reused in different administrative domains.

You can configure, examine, and filter Auto-RP discovery and announcement messages at the administratively scoped boundary. Any Auto-RP group range announcements from the Auto-RP packets that are denied by the boundary ACL are removed. An Auto-RP group range announcement is permitted and passed by the boundary filter only if all addresses in the Auto-RP group range are permitted by the boundary ACL. If any address is not permitted, the entire group range is filtered and removed from the Auto-RP message before the Auto-RP message is forwarded.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Multicast Boundary Filter**, and then click **Add** or **Edit**.

Use the **Add Multicast Boundary Filter** dialog box to add new multicast boundary filters to the device. Use the **Edit Multicast Boundary Filter** dialog box to change existing parameters.

You can configure a multicast boundary for administratively scoped multicast addresses. A multicast boundary restricts multicast data packet flows and enables reuse of the same multicast group address in different administrative domains. When a multicast boundary is defined on an interface, only the multicast traffic permitted by the filter ACL passes through the interface.

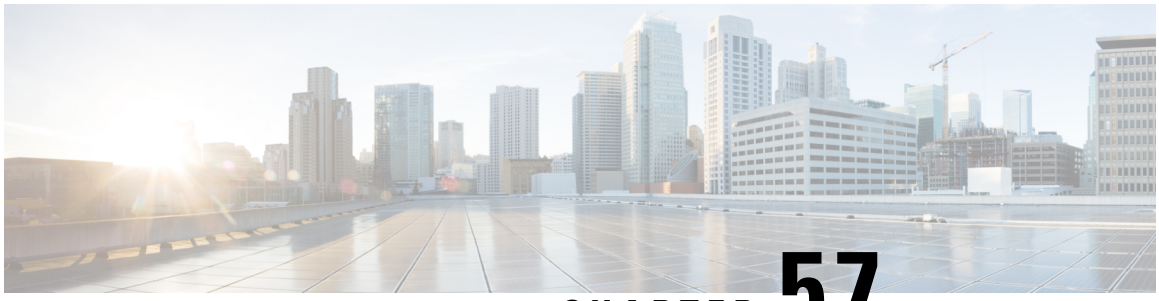
Step 3 From the **Interface** drop-down list, choose the interface for which you are configuring the multicast boundary filter ACL.

Step 4 From the **Standard Access List** drop-down list, choose the standard ACL you want to use, or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object, on page 2953 for the procedure.

Step 5 Check the **Remove any Auto-RP group range announcement from the Auto-RP packets that are denied by the boundary** check box to filter Auto-RP messages from sources denied by the boundary ACL. If this check box is not checked, all Auto-RP messages are passed.

Step 6 Click **OK** to save the multicast boundary filter configuration.

The multicast boundary filter is configured on the selected interface with the specified standard ACL, restricting multicast traffic according to the filter rules and enabling address reuse across administrative domains.



CHAPTER 57

Policy-Based Routing

This chapter describes how to configure Firewall Threat Defense to support policy-based routing (PBR) through Cloud-Delivered Firewall Management Center's policy-based routing page. The following sections describe policy based routing, guidelines for PBR, and configuration for PBR.

- About policy-based routing, on page 1541
- Licenses for policy-based routing, on page 1543
- Best practices for policy-based routing, on page 1543
- Path monitoring metrics for best path route determination, on page 1545
- Configure policy-based routing policy, on page 1548
- Configure policy-based routing, on page 1551
- Configure PBR with path monitoring, on page 1556
- Useful CLIs for monitoring PBR, on page 1558
- Troubleshoot PBR, on page 1560

About policy-based routing

Policy-based routing (PBR) extends traditional routing protocols to provide greater control over traffic flow. It enables routing decisions based on criteria beyond destination IP, such as source and destination ports, protocols, and applications. Furthermore, PBR facilitates secure, application-specific traffic breakout in large-scale network deployments.

Policy-based routing characteristics

In traditional routing, packets are routed based on the destination IP address. Changing the routing of specific traffic in a destination-based routing system is difficult. PBR extends and complements the mechanisms provided by routing protocols.

PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link. With PBR, you can define routing that is based on criteria other than destination network such as source port, destination address, destination port, protocol, applications, or a combination of these objects.

You can use PBR to classify the network traffic based on applications, username, group membership, and security group association. This routing method applies to scenarios involving numerous devices accessing applications and data in large network deployment. In large deployments, network traffic is typically backhauled to a hub as encrypted traffic in a route-based VPN. These topologies often result in issues such as packet

latency, reduced bandwidth, and packet drop. Resolving these issues requires costly, complex deployments and management.

PBR policy enables you to securely break out traffic for specified applications. You can configure PBR policy in the Cloud-Delivered Firewall Management Center user interface to allow the applications to be directly accessed.

Consider a company that has two links between locations: one a high-bandwidth, low-delay expensive link, and the other a low-bandwidth, higher-delay, less-expensive link. While using traditional routing protocols, the higher-bandwidth link gets most, if not all, of the traffic sent across it based on the metric savings obtained by the bandwidth, delay, or both (using EIGRP or OSPF) characteristics of the link. With PBR, you can route higher priority traffic over the high-bandwidth/low-delay link, while sending all other traffic over the low-bandwidth/high-delay link.

These scenarios demonstrate policy-based routing applications:

- **Direct internet access**: In this topology, application traffic from the branch office can be routed directly to the internet instead of through the VPN tunnel connecting to the headquarters. The branch Firewall Threat Defense is configured with an internet exit point. The PBR policy is applied on the ingress interface (*Inside 1*) to identify the traffic based on the applications, user identity (username and group membership), and Security Group Tag (security group association) defined in the ACL. Correspondingly, the traffic is forwarded through the egress interfaces directly to the internet or to the IPsec VPN tunnel.
- **Equal-access and source-sensitive routing**: In this topology, traffic from the HR and Mgmt networks can be configured to go through ISP1 and traffic from Eng network can be configured to go through ISP2. Thus, policy-based routing enables the network administrators to provide equal-access and source-sensitive routing, in this example.
- **Load sharing**: In addition to the dynamic load-sharing capabilities offered by ECMP load balancing, network administrators can now implement policies to distribute traffic among multiple paths based on the traffic characteristics. As an example, in the topology depicted in the Equal-Access Source Sensitive Routing scenario, an administrator can configure policy-based routing to route the traffic from HR network through ISP1 and traffic from Eng network through ISP2 and thus share the load.

Figure 334: Direct internet access scenario

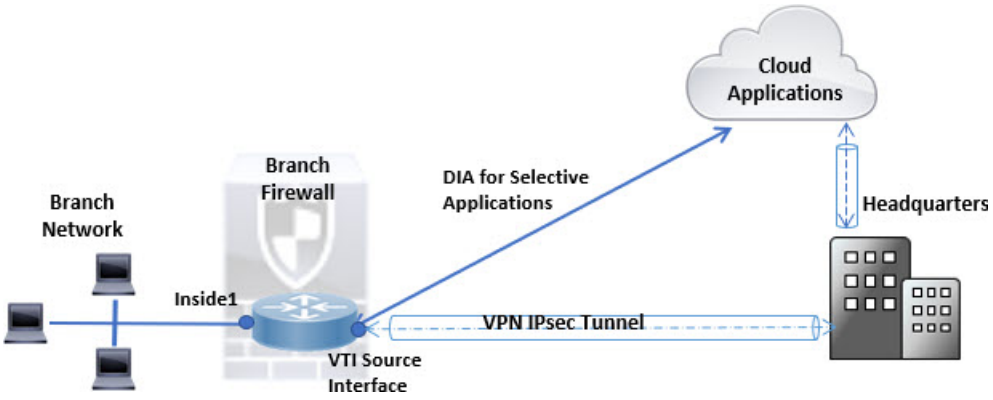
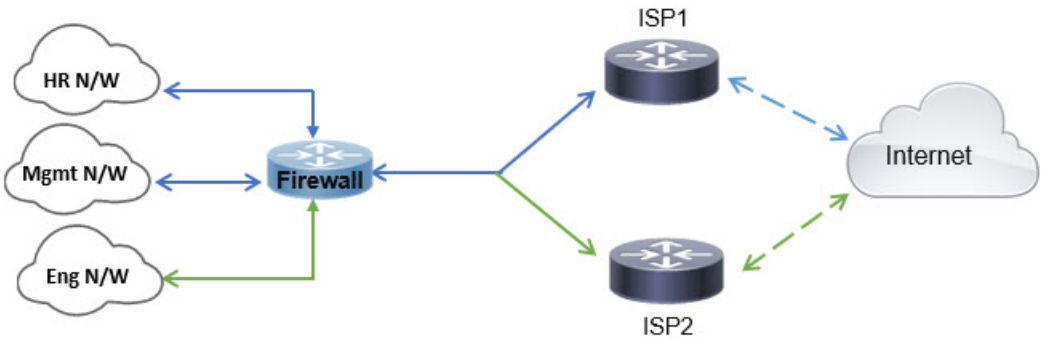


Figure 335: Equal access scenario



Licenses for policy-based routing

Policy-based routing (PBR) has specific licensing requirements depending on the features and services used in your configuration.

- All PBR policy features are supported with the Essentials Firewall Threat Defense smart license.
- Secure Client licensing is required if you are configuring PBR for routing traffic through RA VPN tunnel interface.
- Cisco ISE license is required for ISE usage in PBR policy.

Best practices for policy-based routing

Device guidelines

Ensure proper version compatibility and device-specific configurations when implementing PBR.

- On cluster devices, do not configure application, user identity, and Security Group Tag (SGT) based PBR policies, because they are not supported on cluster devices.

Interface guidelines

Configure appropriate interface types and settings for PBR policies.

- Configure only routed-mode interfaces and non management-only interfaces in the Global virtual router as ingress or egress interfaces for the PBR policy. You cannot configure user-defined virtual router interfaces for the policy.
- The interfaces that you want to define in the policy must have a logical name.
- Configure Static VTIs only as egress interfaces.
- Do not choose Dynamic VTIs for configuring the PBR egress interfaces.

You can apply PBR for managing both IPv4 and IPv6 traffic.

Application-based PBR and DNS configuration

Ensure proper DNS configuration for application-based PBR functionality.

- Application-based PBR uses DNS snooping for application detection. Application detection succeeds only if the DNS requests pass through Firewall Threat Defense in a clear-text format; the DNS traffic is not encrypted.
- You must configure trusted DNS servers for the application detection to succeed.

For more information on configuring DNS servers, refer to Configure DNS server settings, on page 1246.

PBR policies not applied for output route look-up

PBR is an ingress-only feature; The system applies PBR only to the first packet of a new incoming connection, at which time the egress interface for the forward leg of the connection is selected. Note that PBR will not be triggered if the incoming packet belongs to an existing connection, or if NAT is applied and NAT chooses the egress interface.

PBR policies not applied for embryonic traffic

An embryonic connection is where the necessary handshake between source and destination has not been made. When a new internal interface is added and a new VPN policy is created using a unique address pool, PBR is applied to the outside interface matches the source of the new client pool. Thus, PBR sends traffic from the client to the next hop on the new interface. However, PBR does not process return traffic from a host until the new internal interface establishes a connection with the client. Thus, the return traffic from the host to the VPN client, specifically, the VPN client response is dropped as there is no valid route. To prevent the response from being dropped, configure a weighted static route with a higher metric on the internal interface.

HTTP-based path monitoring guidelines

Configure HTTP-based path monitoring with appropriate interface types and version considerations.

- Configure HTTP-based path monitoring only on physical, port-channel, subinterfaces, and static tunnel interfaces. Do not configure it on cluster devices.
- HTTP uses only IPv4 to ping applications. IPv4 metrics are applied for routing and forwarding both IPv4 and IPv6 traffic.
- HTTP-based application monitoring is enabled by default in Cloud-Delivered Firewall Management Center.

Additional guidelines

Follow these additional configuration considerations for effective PBR implementation.

- All existing configuration restrictions and route map limitations remain in effect.
- When you define the ACL for the policy match criteria, you can select multiple predefined applications from a list to form an Access Control Entry (ACE). In Firewall Threat Defense, the predefined applications are stored as Network Service objects and the group of applications as Network Service Groups (NSG). You can create a maximum of 1024 such NSGs. The application or network service group is detected through first-packet classification. Currently, you cannot add to or modify the predefined applications list.
- Unicast Reverse Path Forwarding (uRPF) validates the source IP address of packets received on an interface against the routing table and not against the PBR route map. When uRPF is enabled, packets received on an interface through PBR are dropped as they are without the specific route entry. Therefore, when using PBR, ensure you disable uRPF.

Path monitoring metrics for best path route determination

Path monitoring metrics for best path route determination is a dynamic routing method that uses various monitoring approaches, such as ICMP and HTTP, to collect performance metrics and determine the best path.

Path monitoring methods and operation

PBR uses either static cost or path monitoring (dynamic metrics) to route its traffic.

Path monitoring, when configured on interfaces, derive metrics such as round trip time (RTT), jitter, mean opinion score (MOS), and packet loss per interface. These metrics are used to determine the best path for routing PBR traffic.

ICMP-based path monitoring method

The metrics on the interfaces are collected dynamically using ICMP probe messages to the interface's default gateway or a specified remote peer.

HTTP-based path monitoring method

Path monitoring calculates dynamic metrics for each remote peer associated with an interface. HTTP is preferred over ICMP when monitoring multiple applications and determining the best path through a policy on a branch firewall, because:

- HTTP-ping can derive the performance metrics of the path up to the application layer of the server, where the application is hosted.
- Because the application domain is tracked instead of the IP address, you do not need to change the firewall configuration when the application server IP address changes.



Note You can configure both ICMP and HTTP on the same interface. If the destination in the policy matches to any of the domain IP, the corresponding metrics are used. If the destination does not match to any of the configured domains, the metrics from the ICMP are used by PBR to select the outgoing interface.

Default monitoring timers for metric collection

For metric collection and monitoring, the following timers are used:

- The interface monitor average interval is 30 seconds. This interval indicates the frequency to which the probes average.
- The interface monitor update interval is 30 seconds. This interval indicates the frequency at which the average of the collected values are calculated and made available for PBR to determine the best routing path.
- The interface monitor probe interval by ICMP is one second. This interval indicates the frequency at which an ICMP ping is sent.
- The application monitor probe interval by HTTP is 10 seconds. This interval indicates the frequency at which an HTTP ping is sent. Path monitoring uses the last 30 samples of HTTP ping for calculating the average metrics.

Note You cannot configure or modify the interval for any of these timers.

PBR and path monitoring integration

PBR uses IP-based path monitoring to collect the performance metrics (RTT, jitter, packet-lost, and MOS) of the egress interfaces. PBR uses the metrics to determine the best path (egress interface) for forwarding the traffic. Path monitoring periodically notifies PBR about the monitored interface whose metric got changed. PBR retrieves the latest metric values for the monitored interfaces from the path monitoring database and updates the data path.

Path monitoring functions only with dynamic metrics, and only if the RTT, jitter, packet-lost, or MOS variables are set on the interfaces. Path monitoring does not function with static metrics—interface cost (cost set in interface).

You must enable path monitoring for the interface and configure the monitoring type. The PBR policy page allows you to specify the desired metric for path determination. See Configure policy-based routing policy, on page 1548.

PBR and HTTP-based path monitoring

PBR can be configured to use HTTP-based path monitoring to collect the performance metrics of the application domains and not just one destination IP address. Path monitoring begins only after a DNS entry for a domain is detected; it does not start immediately when HTTP-based application monitoring is configured. After obtaining the resolved IP address for the domain, path monitoring sends an HTTP request and receives a response. If DNS resolves multiple IP addresses for a domain, path monitoring probes and monitors the application using the first resolved IP address. Path monitoring continues until the IP address changes or until HTTP-based monitoring is disabled.

Based on the HTTP request and response durations, path monitoring computes the performance metrics for the application. Path monitoring sends collected metrics to PBR at regular intervals so that PBR can make routing and forwarding decisions for traffic from the configured ingress interface. If traffic arrives before path monitoring sends metrics to PBR, the routing table determines the traffic flow. Once metrics are available, PBR uses them to make routing decisions for subsequent traffic.

Note Based on the Network Service Groups in the match ACL of the policy, you can apply PBR for multiple domains having multiple IP addresses.

The management center associates applications and NSGs with the egress interface only when the PBR configuration meets these criteria:

- The match ACL contains the monitored applications.
- The PBR policy is configured with any one of the interface ordering values (metric type):
  - Minimal Jitter
  - Maximum Mean Opinion Score
  - Minimal Round-Trip Time
  - Minimal Packet Loss

Configure path monitoring settings

Configure path monitoring to collect metrics such as RTT, jitter, MOS, and packet loss on specified interfaces for effective traffic management and routing optimization.

The PBR policy relies on flexible metrics, such as round-trip time (RTT), jitter, mean opinion score (MOS), and packet loss of the interfaces to identify the best routing path for its traffic. Path monitoring collects these metrics on the specified interfaces. On the **Interfaces** page, you can configure path monitoring settings on interfaces to send ICMP probes or HTTP pings to collect metrics.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Path Monitoring** tab.

Step 4 To configure ICMP based monitoring of the interface, click the **Enable IP based Monitoring** check box.

Step 5 From the **Monitoring Type** drop-down list, select the relevant option:

- **Auto**—Sends ICMP probes to the IPv4 default gateway of the interface. If the IPv4 gateway does not exist, path monitoring sends the probes to the IPv6 default gateway of the interface.
- **Peer IPv4**—Sends ICMP probes to the specified peer IPv4 address (next-hop IP) for monitoring. If you select this option, enter the IPv4 address in the **Peer IP To Monitor** field.
- **Peer IPv6**—Sends ICMP probes to the specified peer IPv6 address (next-hop IP) for monitoring. If you select this option, enter the IPv6 address in the **Peer IP To Monitor** field.
- **Auto IPv4**—Sends ICMP probes to the default IPv4 gateway of the interface.
- **Auto IPv6**—Send ICMP probes to the default IPv6 gateway of the interface.

Note

- The Auto options are not available for VTI interfaces. You must specify the peer address.
- Only one next-hop is monitored to a destination. That is, you cannot specify more than one peer address to monitor for an interface.

Step 6 By default, the **Enable HTTP based Application Monitoring** check box is selected. All applications selected for path monitoring in the match ACL of the PBR policy are listed, provided this interface is configured as the egress interface in the policy. To disable the HTTP-based monitoring of the interface, clear the check box.

Step 7 Click **Ok**.

Step 8 To save the settings, click **Save**.

Path monitoring is configured on the interface with the specified settings. The system will collect metrics using ICMP probes or HTTP pings based on your configuration to optimize routing paths.

Add path monitoring dashboard

Add the Path Monitoring Dashboard to enable viewing path monitoring metrics with customizable display options through predefined correlations and widget management.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **System** (⚙) > **Health** > **Monitor**.

Step 2 Select the device, and click **Add new dashboard** (+).

Step 3 Enter a name for the custom dashboard.

Step 4 In the **Metrics** area, click the **Add from predefined correlations** button.

Step 5 From the list, click **Interface - Path metrics**.

By default, all four metrics and an additional metric field are selected to display as widgets in the dashboard.

To exclude any, click **Delete** ( ).

Step 6 Click **Add dashboard**.

Configure policy-based routing policy

This task enables you to configure Policy-Based Routing (PBR) to control traffic routing based on specific criteria rather than relying solely on the routing table.

You can configure the PBR policy on the Policy Based Routing page by specifying the ingress interfaces, match criteria (Extended Access Control List), and egress interfaces.

Before you begin

To use the path monitoring metrics for configuring the traffic forwarding priority over egress interfaces, you must configure the path monitoring settings for the interfaces. See Configure path monitoring settings, on page 1547.

Follow these steps to configure policy-based routing policy:

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Policy Based Routing**.

On the Policy Based Routing page, you can view the configured policy. The grid shows ingress interfaces and the combination of the policy-based route access list and egress interfaces.

Step 4 To configure the policy, click **Add**.

Step 5 In the **Add Policy Based Route** dialog box, select the **Ingress Interface** from the drop-down list.

Note
You can only select interfaces with logical names that belong to a global virtual router from the drop-down list. You cannot configure a VLAN interface with a logical name as the source (ingress) interface.

Step 6 To specify the match criteria and the forward action in the policy, click **Add**.

Step 7 In the **Add Forwarding Actions** dialog box, do the following:

a) From the **Match ACL** drop-down, choose the extended access control list object. You can predefine the ACL object (see Configure an extended ACL object, on page 2949) or click the **Add** (+) icon to create the object. In the **New Extended Access List Object** box, enter a name, click **Add** to open the **Add Extended Access List Entry** dialog box, where you can define the network, port, user identity, SGT, or application match criteria for the PBR policy.

Note
You can either have destination address or application/user identity/SGT defined in an ACE.

To selectively apply PBR on the incoming interface, you can define *Block* criteria in the ACE. When the traffic matches the block rule of the ACE, the traffic is forwarded to the egress interface based on the routing table.

b) From the **Send To** drop-down list:

- To select the configured interfaces, choose **Egress Interfaces**.
- To specify the IPv4/IPv6 next hop addresses, choose **IP Address**. Proceed to Step 7.e, on page 1550

c) If you have selected **Egress Interfaces**, from the **Interface Ordering** drop-down, choose the relevant option:

- By **Interface Priority**—The traffic is forwarded based on the priority of the interfaces. Traffic is routed to the interface with the least priority value first. When the interface is not available, the traffic is then forwarded to the interface with the next lowest priority value. For example, let us assume that interface1, interface2, and interface3 are configured with priority values *0*,*1*, and *2*, respectively. The traffic is forwarded to interface1. If interface1 becomes unavailable, the traffic is then forwarded to interface2.

  Note
  To configure the priority for the interfaces, click **Configure Interface Priority** on the Policy Based Routing page. In the dialog box, provide the priority number against the interfaces, and then click **Save**. You can also configure the priority for an interface in the Interface Settings.

  When the priority value is the same for all the interfaces, the traffic is balanced among the interfaces. By default, the interface priority for PBR is set to 0. Interfaces with the lowest priority value are preferred for traffic forwarding.

- By **Order**—The traffic is forwarded based on the sequence of the interfaces specified here. For example, let us assume that interface1, interface2, and interface3 are selected in this order, interface2, interface3, interface1. The traffic is forwarded first to interface2, then to interface3, regardless of their priority values.
- By **Minimal Jitter**—The traffic is forwarded to the interface that has the lowest jitter value. You need to enable Path Monitoring on the interfaces for PBR to obtain the jitter values.

- By **Maximum Mean Opinion Score**—The traffic is forwarded to the interface that has the maximum mean opinion score (MOS). You need to enable Path Monitoring on the interfaces for PBR to obtain the MOS values.
- By **Minimal Round Trip Time**—The traffic is forwarded to the interface that has the minimal round trip time (RTT). You need to enable Path Monitoring on the interfaces for PBR to obtain the RTT values.
- By **Minimal Packet Loss**—The traffic is forwarded to the interface that has the minimal packet loss. You need to enable Path Monitoring on the interfaces for PBR to obtain the packet loss values.

d) In the **Available Interfaces** box, all the interfaces with their priority values are listed. From the list of interfaces, click the **Add** (+) button to add to the selected egress interfaces. Proceed to Step 8.

Note
The selected interface route must be present in the routing table.

e) If you have selected **IP Address**, enter the IP addresses separated by commas in the **IPv4 Addresses** or **IPv6 Addresses** fields.

Note
When multiple next-hop IP addresses are provided, the traffic is forwarded as per the sequence of the specified IP addresses until a valid routable next-hop IP address is found. The configured next-hops should be directly connected.

f) From the **Don't Fragment** drop-down list, select Yes, No, or None. If the DF (Don't Fragment) flag is set to *Yes*, the intermediate routers never perform fragmentation of a packet.To specify the current interface as the default for forwarding, check the **Default Interface** check box.

g) The **IPv4 Settings** and **IPv6 Settings** tabs enable you to specify the recursive and default settings:

Note
For a route-map, you can only specify either IPv4 or IPv6 next-hop settings.

- **Recursive**—The route map configuration is applied only when the specified next-hop address and the default next-hop address are found on a directly connected subnet. However, you could use the recursive option, where the next-hop address need not be directly connected. Here, a recursive lookup is performed on the next-hop address, and matching traffic is forwarded to the next-hop used by that route entry according to the current routing path of the router.
- **Default**—If the normal route lookup fails to match traffic, the traffic is forwarded to this specified next-hop IP address.

h) Check the **Peer Address** check box to use the next-hop address as the peer address.

Note
You cannot configure a route map with both default next-hop address and peer address.

i) For IPv4 settings, you can check whether the next IPv4 hops of a route map are available under **Verify Availability**—click the **Add** (+) button and add the next-hop IP address entries:

- **IP Address**—Enter the next hop IP address.
- **Sequence**—Entries are assessed in order using the sequence number. Ensure that no duplicate sequence numbers are entered. The valid range is 1 to 65535.

- **Track**—Enter a valid ID. The valid range is 1 to 255.

Step 8 To save the policy, click **Save** and **Deploy**.

The Firewall Threat Defense uses ACLs to match traffic and perform routing actions on the traffic. Typically, you configure a route map that specifies an ACL against which traffic is matched, and then you specify one or more actions for that traffic. With the use of path monitoring, PBR can now select the best egress interface for routing the traffic. Finally, you associate the route map with an interface on which you want to apply PBR on all incoming traffic.

Configure policy-based routing

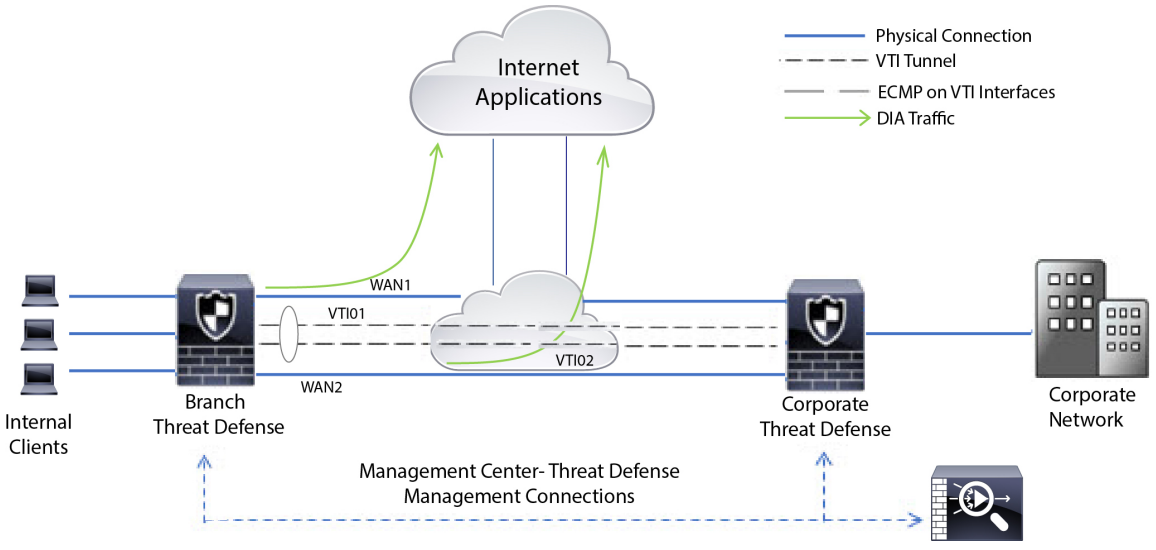
Configure PBR to enable direct internet access for specific web-based applications, optimize traffic flow, and achieve load balancing between WAN and VTI interfaces.

Consider a typical corporate network scenario where all the branch network traffic passes through a route-based VPN of the corporate network and diverges to the extranet, when required. Accessing web-based applications through the corporate network to support daily operations may increase network size and maintenance costs. This example illustrates how to configure PBR for direct internet access.

The following figure depicts the topology of a corporate network. The branch network is connected to the corporate network through a route-based VPN. Traditionally, the corporate Firewall Threat Defense is configured to handle both the internal and external traffic of the branch office. With the PBR policy, the branch Firewall Threat Defense is configured with a policy that routes specific traffic to the WAN network instead of the virtual tunnels. The rest of the traffic flows through the route-based VPN, as usual.

This example also illustrates how to configure the WAN and VTI interfaces with ECMP zones to achieve load balancing.

Figure 336: Configuring policy-based routing on Branch Firewall Threat Defense in Cloud-Delivered Firewall Management Center



Before you begin

This example assumes that you have already configured WAN and VTI interfaces for the branch Firewall Threat Defense in Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Configure policy-based routing for the branch Firewall Threat Defense, select the ingress interfaces:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Choose **Routing** > **Policy Based Routing**, and on the **Policy Based Routing** page, click **Add**.

c) In the **Add Policy Based Route** dialog box, select the interfaces (say, *Inside 1*, and *Inside 2*) from the **Ingress Interface** drop-down list.

Step 2 Specify the match criteria:

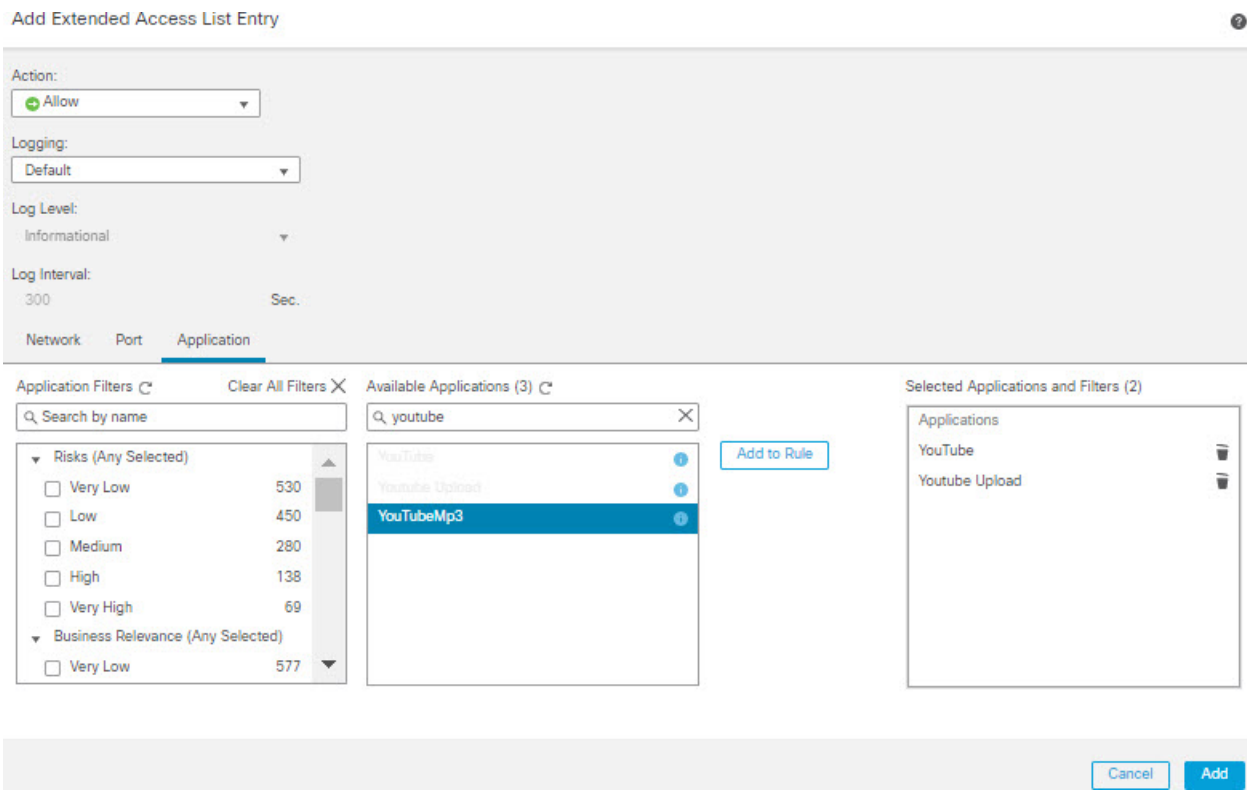
a) Click **Add**.

b) To define the match criteria, click the **Add** (+) button.

c) In **New Extended Access List Object**, enter the name for the ACL (say, *DIA-FTD-Branch*), and click **Add**.

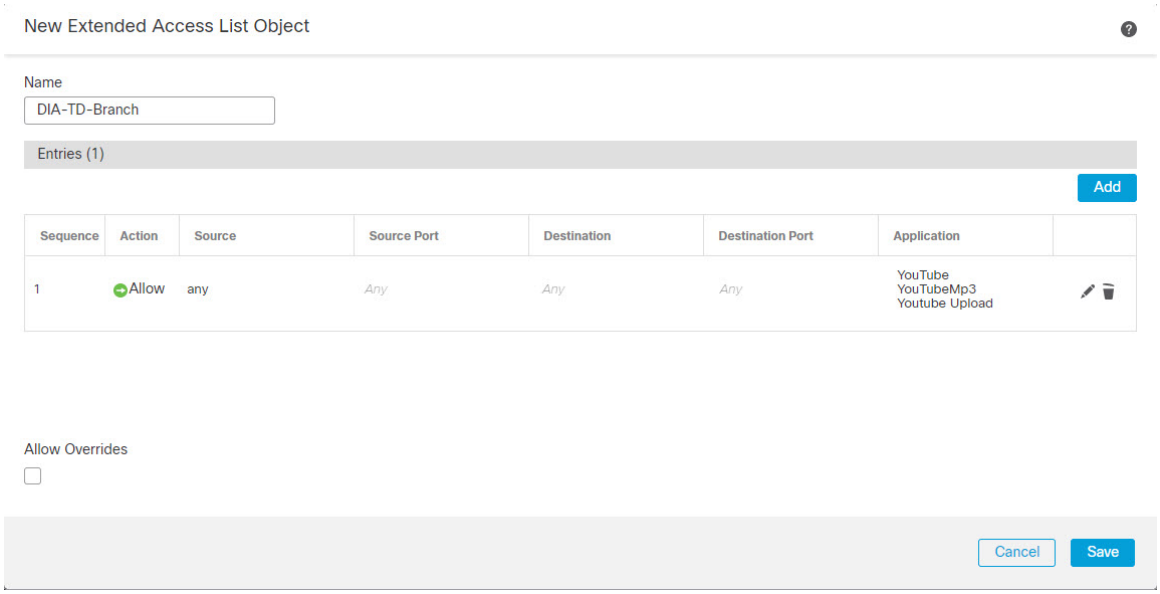
d) In the **Add Extended Access List Entry** dialog box, choose the required web-based applications from the **Application** tab:

Figure 337: Applications tab



On the Firewall Threat Defense, the application group in an ACL is configured as a network service group and each of the applications as a network service object.

Figure 338: Extended ACL



e) Click **Save**.

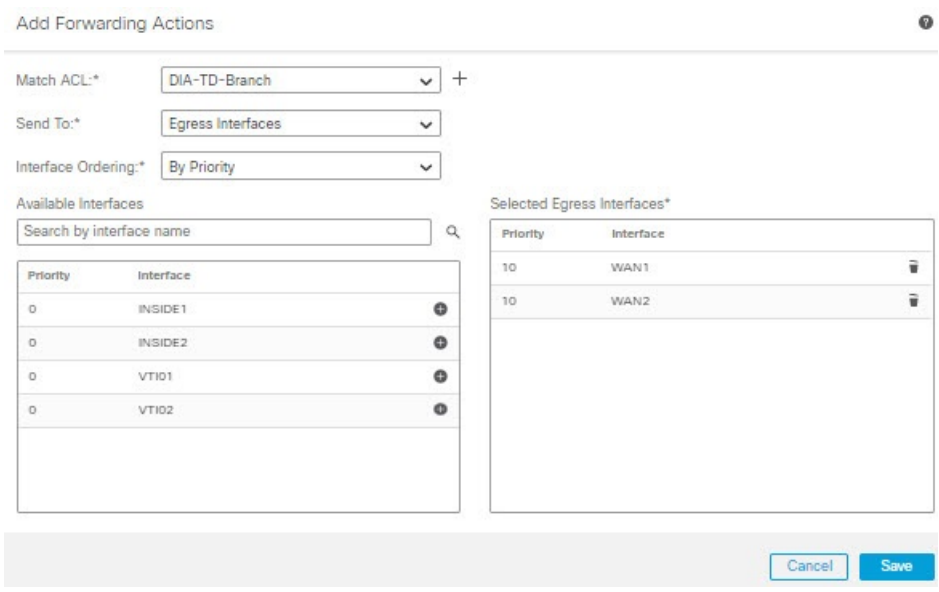
f) Select *DIA-FTD-Branch* from the **Match ACL** drop-down list.

Step 3 Specify the egress interfaces:

a) From the **Send To** and **Interface Ordering** drop-down lists, choose Egress Interfaces, and Interface Priority respectively.

b) Under **Available Interfaces**, click the ⊕ button against the respective interface names to add *WAN1* and *WAN2*:

Figure 339: Configuring policy-based routing



c) Click **Save**.

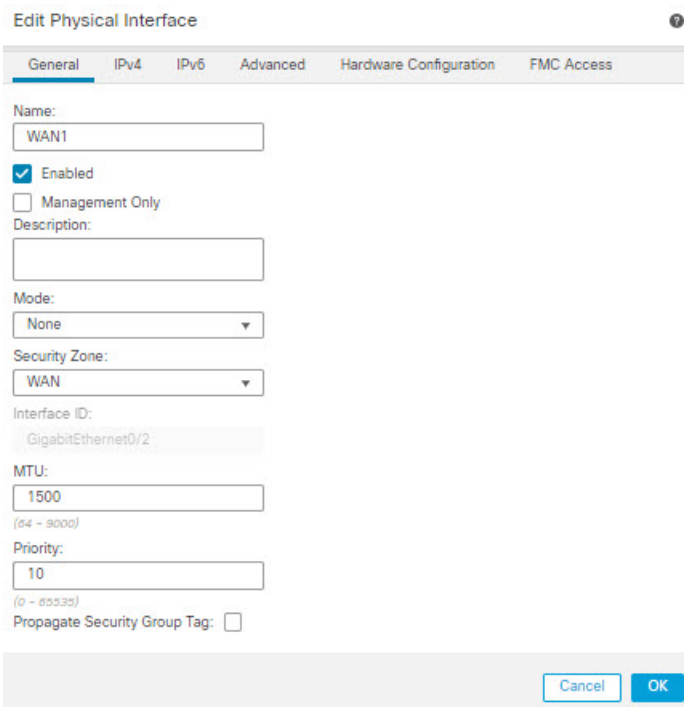
Step 4 Configure interface priority:

You can set the priority value for the interfaces either in the **Edit Physical Interface** page, or in the **Policy Based Routing** page (**Configure Interface Priority**). In this example, the Edit Physical Interface method is described.

a) Choose **Devices** > **Device Management**, and edit the branch Firewall Threat Defense.

b) Set the priority for the interfaces. Click **Edit** against the interface and enter the priority value:

Figure 340: Setting interface priority



c) Click **Ok** and **Save**.

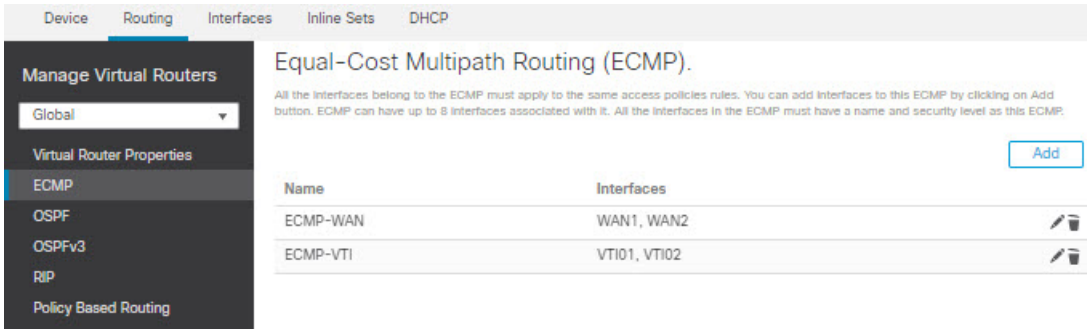
Step 5 Create ECMP zones for load balancing:

a) In the **Routing** page, click **ECMP**.

b) To associate interfaces to the ECMP zone, click **Add**.

c) Select *WAN1* and *WAN 2* and create an ECMP zone—*ECMP-WAN*. Similarly, add *VTI01* and *VTI02* and create an ECMP zone—*ECMP-VTI*:

Figure 341: Associating interfaces with ECMP zone

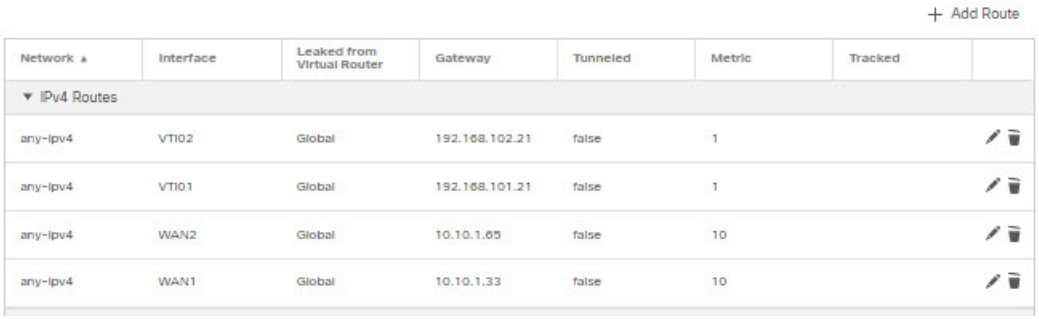


Step 6 Configure static routes for the zone interfaces for load balancing:

a) In the **Routing** page, click **Static Route**.

b) Click **Add** and specify the static routes for *WAN1*, *WAN2*, *VTI01*, and *VTI02*. Ensure that you specify the same metric value for the interfaces belonging to the same ECMP zones (Step 5):

Figure 342: Configuring static routes for ECMP zone interfaces



Note
Ensure that the zone interfaces have the same destination address and metric, but different gateway addresses.

Step 7 Configure trusted DNS on the WAN objects of the branch Firewall Threat Defense to ensure secure flow of traffic to the internet:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Platform Settings**, and create a DNS policy on the branch Firewall Threat Defense.

b) To specify the trusted DNS, edit the policy and click **DNS**.

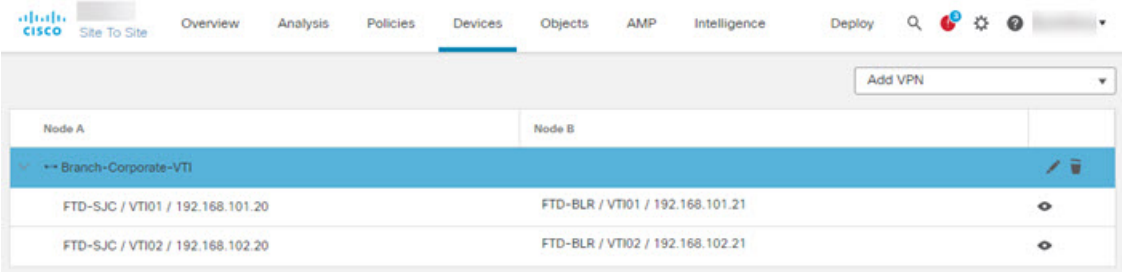
c) To specify the DNS servers for the DNS resolution to be used by WAN objects, in the **DNS Settings** tab, provide the DNS server group details and select WAN from the interface objects.

d) Use the **Trusted DNS Servers** tab to provide specific DNS servers that you trust for the DNS resolution.

Step 8 Click **Save**, and then click **Deploy**.

Any *YouTube* related access requests from the branch inside network *INSIDE1* or *INSIDE2* are routed to *WAN1* or *WAN2* as they would match the *DIA-FTD-Branch* ACL. Any other request, say *google.com*, are routed through *VTI01* or *VTI02* as configured in the Site to Site VPN Settings.

Figure 343: Site to site VPN settings



With the ECMP configured, the network traffic is seamlessly balanced.

Configure PBR with path monitoring

Configure PBR with path monitoring to optimize application traffic routing by dynamically selecting the best egress interface based on performance metrics such as jitter, round-trip time, and packet loss.

This example describes how to configure PBR with path monitoring for the following applications with flexible metrics:

- Audio or video sensitive applications (example, WebEx Meetings) with Jitter.
- Cloud-based application (example, Office365) with RTT.
- Network-based access control (with a specific source and destination) with packet loss.

Before you begin

1. This example assumes that you are aware of the basic configuration steps for PBR.
2. You have configured ingress and egress interfaces with logical names. In this example, the ingress interface is named *Inside1*, and egress interfaces are named *ISP01*, *ISP02*, and *ISP03*.

Follow these steps to configure PBR with path monitoring:

Procedure

Step 1 Path monitoring configuration on interfaces *ISP01*, *ISP02*, and *ISP03*:

For the metrics collection on the egress interfaces, you must enable and configure path monitoring on them.

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit Firewall Threat Defense.
b) Under the **Interfaces** tab, edit the interface (in our example, *ISP01*)
c) Click the **Path Monitoring** tab, select the **Enable Path Monitoring** check box, and then specify the monitoring type (see Configure path monitoring settings, on page 1547).
d) Click **Ok** and **Save**.
e) Repeat these steps to configure the path monitoring settings for *ISP02* and *ISP03*.

Step 2 Configure policy-based routing for a branch in an organization Firewall Threat Defense, select the ingress interfaces:

a) In the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Policy Based Routing**, and on the **Policy Based Routing** page, click **Add**.
c) In the **Add Policy Based Route** dialog box, select *Inside 1* from the **Ingress Interface** drop-down list.

Step 3 Specify the match criteria:

a) Click **Add**.
b) To define the match criteria, click the **Add** (+) button.
c) In **New Extended Access List Object**, enter the name for the ACL (example, *PBR-WebEx*), and click **Add**.
d) In the **Add Extended Access List Entry** dialog box, choose the necessary web-based applications, such as WebEx Meetings) from the **Application** tab.

Remember
On Firewall Threat Defense, the application group in an ACL is configured as a network service group and each of the applications as a network service object.

e) Click **Save**.
f) Select *PBR-WebEx* from the **Match ACL** drop-down list.

Step 4 Specify the egress interfaces:

a) From the **Send To** drop-down list, choose Egress Interfaces.
b) From the **Interface Ordering** drop-down list, choose By Minimal Jitter.
c) Under **Available Interfaces**, click the **Right Arrow** (>)button against the respective interface names to add *ISP01*, *ISP02*, and *ISP03*.
d) Click **Save**.

Step 5 Repeat Step 2 and Step 3 to create PBRs for the *Inside1* interface to route Office365 and network-based access control traffic:

a) Create a match criteria object, example *PBR-Office365*, and select the Office365 application from the **Application** tab.
b) From the **Interface Ordering** drop-down list, choose By Minimal Round Trip Time.
c) Specify the egress interfaces *ISP01*, *ISP02*, and *ISP03*, and click **Save**.
d) Now, create a match criteria object, example *PBR-networks*, and specify the source and destination interface in the **Network** tab.
e) From the **Interface Ordering** drop-down list, choose By Minimal Packet Loss.
f) Specify the egress interfaces *ISP01*, *ISP02*, and *ISP03*, and click **Save**.

Step 6 Click **Save**, and then click **Deploy**.

Step 7 To view path monitoring metrics, in the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management**, and from **More** (⋮) click **Health Monitor**. To view the metric details for the interfaces of the device, you must add the path metrics dashboard. For details, see Add path monitoring dashboard, on page 1548.

The WebEx, Office365, and networks-based ACL traffic are forwarded through the best route derived from the metrics value collected on *ISP01*, *ISP02*, and *ISP03*.

Useful CLIs for monitoring PBR

Run the monitoring commands described in this topic from the Firewall Threat Defense device CLI.

Interface configurations

To view the interface configurations of the device, run the `show run interface` command:

```
> show run interface
!
interface Ethernetl/l
 description Outside ispl handoff
 nameif outside1
 security-level 0
 zone-member ECMP-WAN
 ip address dhcp setroute
 policy-route cost 10
 policy-route path-monitoring 8.8.8.8
 policy-route path-monitoring object-group network-service FMC_NSG_4295470581 policy-route
 path-monitoring object-group network-service FMC_NSG_4295470600
!
interface Ethernet1/2
 description Outside isp2 handoff
 nameif outside2
 security-level 0
 zone-member ECMP-WAN
 ip address 192.133.243.240 255.255.255.192
 policy-route cost 20
 policy-route path-monitoring 8.8.8.8
 policy-route path-monitoring object-group network-service FMC_NSG_4295470581 policy-route
 path-monitoring object-group network-service FMC_NSG_4295470600
!
```

DNS configurations

Application-based routing only uses trusted DNS servers to resolve domains. To view the DNS configurations of the device, run the `show run DNS` command:

```
> show run dns
DNS server-group DefaultDNS
dns trusted-source 10.100.0.5
dns trusted-source 10.200.0.5
```

Route map configurations

When you configure PBR on the device, the Management Center auto-generates the route-map and applies it to the specified ingress interface. To view the route maps of the device, run the `show run route-map` command:

```
> show run route-map
!
route-map FMC_VPN_CONNECTED_DIST_RMAP_1000 permit 10
 match interface inside-employee
 set community 1000
!
route-map FMC GENERATED PBR 1729024850865 permit 5
 match ip address Cloud-storage-apps-acl
  set adaptive-intertace cost outside1 outside2
```

```
!
route-map FMC_GENERATED_PBR_1729024850865 permit 10
match ip address Social-media-apps-acl
set adaptive-interface rtt outside1 outside2

!
route-map FMC_GENERATED_PBR_1729024850865 permit 15
match ip address Conferencing-apps-acl
set adaptive-interface jitter outside1 outside2

!
route-map FMC_GENERATED_PBR_1729024850865 permit 20
match ip address Corp-internal-apps-acl
set adaptive-interface cost outside1_static_vti_1 outside2_static_vti_4
```

Access lists and network service groups configurations

The route-map applied to the ingress interface can reference an extended access control list. To view the details of an access list for PBR, run the `show run access list <access list_name>` command:

```
> show run access-list Cloud-storage-apps-acl
access-list Cloud-storage-apps-acl extended permit ip any object-group-network-service
FMC_NSG_4295470562
```

The network-service objects and object-groups are configured in extended access control lists and referenced in policy-based routing route maps and access control groups. To view the NSG configurations, run the `show object-group network-service <network-service-groups-name>` command. The *network-service-groups-name* is derived from the above show command for an access list.

```
> show object-group network-service FMC_NSG_4295470562
 object-group network-servire FMC_NSG_4295470562 (id=@xfdff0000)
 network-service-member "Box" dynamic
 description File storage and transfer site.
 app-id 1326
 domain box.com (bid=436735707) ip (hitcnt=0)
 domain boxcloud.com (bid=436924171) ip (hitcnt=0)
 domain box.net (bid=437080553) ip (hitcnt=0)
 domain box.org (bid=437174273) ip (hitcnt=0)
 domain boxcdn.net (bid=437272231) ip (hitcnt=0)
 domain boxrelay.com (bid=437481703) ip (hitcnt=0)
 domain boxenterprise.net (bid=437626005) ip {hitcnt=0)
 domain boxinvestorrelations.com (bid=437672765) ip (hitcnt=0)
 domain segment-box.com (bid=437886771) ip (hitcnt=0)
 domain box-corp.com (bid=437924995) ip (hitcnt=0)
 domain boxcn.net (bid=438072833) ip (hitcnt=0)
network-service-member "Dropbox" dynamic
 description Cloud based tile storage.
 app-id 125
 domain dropbox.com (bid=24259639) ip (hitcnt=0)
 domain cfl.dropboxstatic.com (bid=24495525) ip (hitcnt=0)
 domain dl.dropboxusercontent.com (bid=24596237) ip (hitcnt=0)
 domain dropboxapi.com (bid=24694467) ip (hitcnt=0)
 domain dropboxbusiness.com (bid=24859859) ip (hitcnt=0)
 domain dropboxcaptcha.com (bid=25008145) ip {hitcnt=0)
 domain dropbox-dns.com (bid=25087753) ip (hitcnt=0)
 domain dropboxer.net (bid=25236751) ip (hitcnt=0)
 domain dropboxusercontent.com (bid=25324335) ip (hitcnt=0)
 domain getdropbox.com (bid=25437501) ip (hitcnt=0)
 domain cloudon.com (bid=25580229) ip (hitcnt=0)
```

Path monitoring configurations

To view the path monitoring metrics collected on the egress interface, run the `show path-monitor` command:

```
> show path-monitor
Interface: outside2 (Ethernetl/2)
Remote peer: 8.8.8.8
    Remote peer reachable: Yes
    RTT average: 9138 microsecondes) Jitter: 1093 microsecond(s)
    Packet loss: 0% MOS: 4.39
    Last updated: 12 second(s) ago
Interface: outside2 (Ethernetl/2)
Remote NSG: FMC_NSG_4295470581
    Network Service: Facebook Domain name: fbsbx.com Remote peer reachable: Yes
    RTT average: 17460 microsecond(s) Jitter: 911 microseconde)
    Packet loss: 0%
    MOS: 4.39
    Last updated: 12 second(s) ago

    Network Service: Facebook
    Domain name: facebook.net
    Remote peer reachable: Yes
    RTT average: 17444 microsecondes)
    Jitter: 836 microsecondes)
    Packet loss: 0%
    MOS: 4.39
    Last updated: 12 second(s) ago

    Network Service: Instagram
    Domain name: instagram.com Remote peer reachable: Yes
    RTT average: 17576 microsecondes)
    Jitter: 429 microsecondes)
    Packet loss: 0%
    MOS: 4.39
    Last updated: 12 secondes) ago

Interface: outside2 (Ethernetl/2)
Remote NSG: FMC_NSG_4295470600
    Network Service: WebEx
    Domain name: webex.com Remote peer reachable: Yes RTT average: 18537 microsecond(s)
Jitter: 318 microseconde)
    Packet loss: 0%
    MOS: 4.39
    Last updated: 12 second(s) ago
    Network Service: Zoom Domain name: zoom.com Remote peer reachable: Yes
    RTT average: 98196 microsecond(s) Jitter: 4120 microseconde)
    Packet loss: 0%
    MOS: 4.34
    Last updated: 12 second(s) ago
```

Troubleshoot PBR

This task helps you debug PBR configuration when packets get dropped by verifying routes, monitoring metrics, and validating interface selection.

When PBR is not working as intended, you need to systematically verify the configuration components including route maps, path monitoring, and interface selection to identify the root cause of packet drops.

Before you begin

Follow these steps to troubleshoot PBR configuration when packets get dropped:

Procedure

Step 1 Check that all essential routes for recursive resolution are present by using the show route or show ipv6 route commands in the appropriate tables.

Unless the route map for the interfaces participating in PBR is updated with correct routes, PBR will not work as intended.

Step 2 Display the remote address that path-monitoring uses to monitor metrics by using the **show running-config interface** command.

Example:

```
interface GigabitEthernet0/0
nameif outside_0
security-level 0
zone-member ecmp-zone
ip address 20.0.0.3 255.255.255.0 -> This is egress interface "show running-config" output,
 the monitored address and cost metric value is determined in this output.
policy-route cost 1
policy-route path-monitoring
20.0.0.4
!
int GigabitEthernet 0/3
!
interface GigabitEthernet0/3
nameif outside_3
security-level 0
ip address 11.1.1.2 255.255.255.0 -> This is ingress interface "show running-config" output,
 the specific route-map will be used by PBR to determine the next route.
policy-route route-map rtt-tes
```

This command shows the monitored address and cost metric values for path monitoring by sending ICMP packets or HTTP pings.

Step 3 Examine the **show path-monitoring** and **show run route-map** outputs for packet loss.

Example:

```
ciscoasa(config)# show path-monitoring
Interface: outside_0  -> The remote address used for ICMP monitoring
Remote peer: 20.0.0.4
Version: 6223
Remote peer reachable: Yes -> If this value turns "No", then the ICMPv4/v6 packet is not
reaching the required remote address.
RTT average: 1920 microsecond(s)
Jitter: 394 microsecond(s)
Packet loss: 0%
MOS: 4.40
Last updated: 17 second(s) ago -> The data should be updated by path monitoring after every
 30 seconds. The 'show route-map' would show the
updated metric values.
Interface: outside_2
Remote peer: 40.0.0.4
Version: 6223
Remote peer reachable: Yes
RTT average: 1935 microsecond(s)
Jitter: 433 microsecond(s)
Packet loss: 0%
MOS: 4.40
Last updated: 17 second(s) ago
```

Example:

```
ciscoasa(config)# show route-map
route-map rtt-test, permit, sequence 10
Match clauses:
Set clauses:
adaptive-interface rtt outside_0 (1920) outside_2 (1935) outside_1 (1971) -> Displays the
metric type (rtt) that is used by the policy route to select the adequate interface to send
 the packet. The interface list where cost of each interface is given. As the metric type
is "rtt" and considering the minimum rtt value, the "outside_0" interface route will be
selected by PBR.
route-map mos-test, permit, sequence 10
Match clauses:
Set clauses:
adaptive-interface mos outside_0 (378) outside_1 (390) outside_2 (440) - > As the metric
type is "mos", considering the maximum value of mos, the "outside_2" interface will be
selected by PBR.
```

The metric types (lost, rtt, jitter, cost) should select the interface with the minimum metric value for routing.

The metric type "mos" should select the interface with the maximum metric value for routing.

Step 4 Use **packet-tracer** command to validate the interface selection based on the metric type defined in the PBR.

Example:

```
packet-tracer input <interface> icmp <src ip address> 8 0 <dst ip
address> detailed
Phase: 3
Type: PBR-LOOKUP
Subtype: policy-route
Result: ALLOW
Elapsed time: 60656 ns
Config:
route-map rtt-test permit 10
match ip address allow_101_1_1_2
set adaptive-interface rtt outside_0 outside_2
Additional Information:
Matched route-map rtt-test, sequence 10, permit
Found next-hop 40.0.0.4 using egress ifc outside_2 -> PBR selects the adequate interface
from adaptiveinterface list given in "rtt-test" route-map.
```

Step 5 Use the **packet-tracer** command similar to the **debug policy-route** command.

When PBR successfully selects the route, the packet tracer output looks like this:

```
pbr: policy based route lookup called for 101.1.1.1/0 to 101.1.1.2/0 proto 1 sub_proto 8
received on
interface outside_3, NSGs, nsg_id=none
pbr: First matching rule from ACL(-1)
pbr: route map rtt-test, sequence 10, permit; proceed with policy routing
pbr: policy based routing applied; egress_ifc = outside_2 : next_hop = 20.0.0.4
```

When PBR is not able to find the adequate route, it falls back to normal route lookup, and the packet tracer output looks like this:

```
pbr: policy based route lookup called for 100.1.1.1/0 to 100.1.1.2/0 proto 1 sub_proto 8
received on interface outside_3, NSGs, nsg_id=none
pbr: First matching rule from ACL(-1)
pbr: route map mos-test, sequence 10, permit; proceed with policy routing
pbr: no route to 100.1.1.2 on adaptive-interface outside_2
pbr: no route to 100.1.1.2 on adaptive-interface outside_1
pbr: no route to 100.1.1.2 on adaptive-interface outside_0
pbr: policy based routing could not be applied; proceeding with normal route lookup
```

When a monitored remote address goes down, and path monitoring marks **Remote peer reachable** as **No** for that address, the PBR displays the a log that excludes the interface from adaptive-interface list.

```
pbr: policy based route lookup called for 100.1.1.1/0 to 101.1.1.2/0 proto 1 sub_proto 8
received on
interface outside_3, NSGs, nsg_id=none
pbr: First matching rule from ACL(-1)
pbr: route map rtt-test, sequence 10, permit; proceed with policy routing
pbr: Path Monitoring Ifc Down : adaptive-interface outside_1 Excluded from PBR routing
pbr: policy based routing applied; egress_ifc = outside_2 : next_hop = 40.0.0.4
```

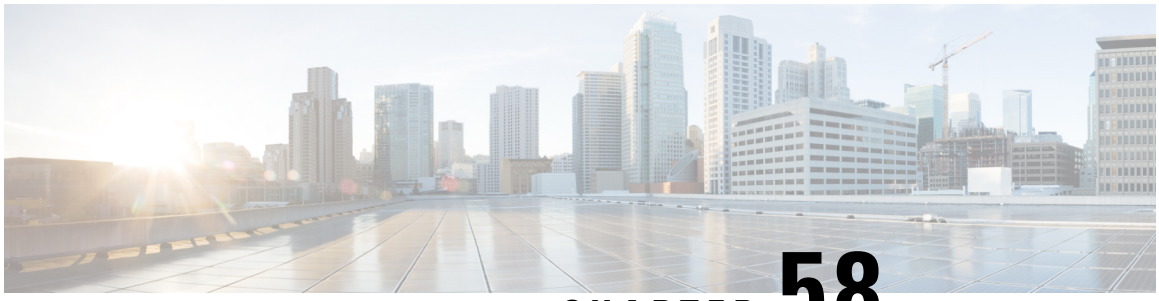
Note
The interface becomes eligible for adaptive PBR routing when the interface is reported as reachable in the path monitoring module.



PART XIV

Network Policies

- Network Address Translation, on page 1567
- Quality of Service, on page 1693



CHAPTER 58

Network Address Translation

The following topics explain Network Address Translation (NAT) and how to configure it on Firewall Threat Defense devices.

- NAT functions, on page 1567
- NAT basics, on page 1568
- Requirements and prerequisites for NAT policies, on page 1577
- Guidelines for NAT, on page 1577
- Manage NAT policies, on page 1583
- Configure NAT for threat defense, on page 1585
- Translating IPv6 networks, on page 1627
- Monitoring NAT, on page 1640
- Examples for NAT, on page 1641

NAT functions

NAT enables private IP networks to connect to the Internet by replacing private IP addresses with public IP addresses. It conserves public addresses by advertising at minimum only one public address for the entire network to the outside world.

Private IP address ranges

Most devices within an IP network are assigned private IP addresses, which are not routable outside the company network due to the shortage of public IPv4 addresses. RFC 1918 defines the private IP address ranges that should not be advertised externally:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

One of the primary functions of NAT is to enable private IP networks to access the Internet by translating private IP addresses into legal, routable public addresses. This approach conserves public IP addresses since NAT can be configured to advertise only a single public address for the entire network.

Other functions of NAT include:

- Security: Hides internal IP addresses, discouraging direct attacks.

- IP routing solutions: Handles overlapping IP addresses without conflict.
- Flexibility: Allows internal IP addressing schemes to change without affecting externally visible public addresses.
- Translating between IPv4 and IPv6 (Routed mode only): Supports translating between IPv4 and IPv6, enabling interoperability between different network types.



Note NAT is optional. If NAT is not configured for specific traffic, that traffic will not be translated but will still be subject to normal security policies.

NAT basics

The following topics explain some of the basics of NAT.

NAT terminology

This document uses the following terminology:

- Real address/host/network/interface—The real address is the address that is defined on the host, before it is translated. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the inside network would be the "real" network. Note that you can translate any network connected to the device, not just an inside network. Therefore if you configure NAT to translate outside addresses, "real" can refer to the outside network when it accesses the inside network.
- Mapped address/host/network/interface—The mapped address is the address that the real address is translated to. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the outside network would be the "mapped" network.



Note During address translation, IP addresses configured for the device interfaces are not translated.

- Bidirectional initiation—Static NAT allows connections to be initiated *bidirectionally*, meaning both to the host and from the host.
- Source and destination NAT—For any given packet, both the source and destination IP addresses are compared to the NAT rules, and one or both can be translated/untranslated. For static NAT, the rule is bidirectional, so be aware that "source" and "destination" are used in commands and descriptions throughout this guide even though a given connection might originate at the "destination" address.

NAT types

You can implement NAT using the following methods:

- Dynamic NAT—A group of real IP addresses are mapped to a (usually smaller) group of mapped IP addresses, on a first come, first served basis. Only the real host can initiate traffic. See Dynamic NAT, on page 1593.
- Dynamic Port Address Translation (PAT)—A group of real IP addresses are mapped to a single IP address using a unique source port of that IP address. See Dynamic PAT, on page 1598.
- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See Static NAT, on page 1608.
- Identity NAT—A real address is statically translated to itself, essentially bypassing NAT. You might want to configure NAT this way when you want to translate a large group of addresses, but then want to exempt a smaller subset of addresses. See Identity NAT, on page 1617.

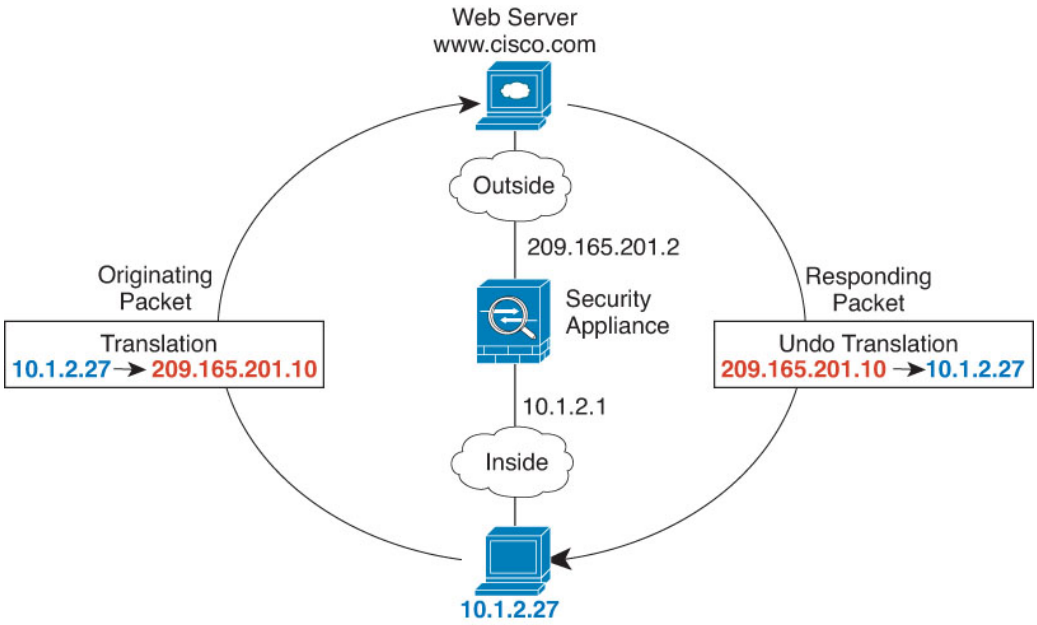
NAT in routed and transparent mode

You can configure NAT in both routed and transparent firewall mode. You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes. The following sections describe typical usage for each firewall mode.

NAT in routed mode

The following figure shows a typical NAT example in routed mode, with a private network on the inside.

Figure 344: NAT Example: Routed Mode



1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is translated to a mapped address, 209.165.201.10.
2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the Firewall Threat Defense receives the packet because the Firewall Threat Defense performs proxy ARP to claim the packet.

3. The Firewall Threat Defense then changes the translation of the mapped address, 209.165.201.10, back to the real address, 10.1.2.27, before sending it to the host.

NAT in transparent mode or within a bridge group

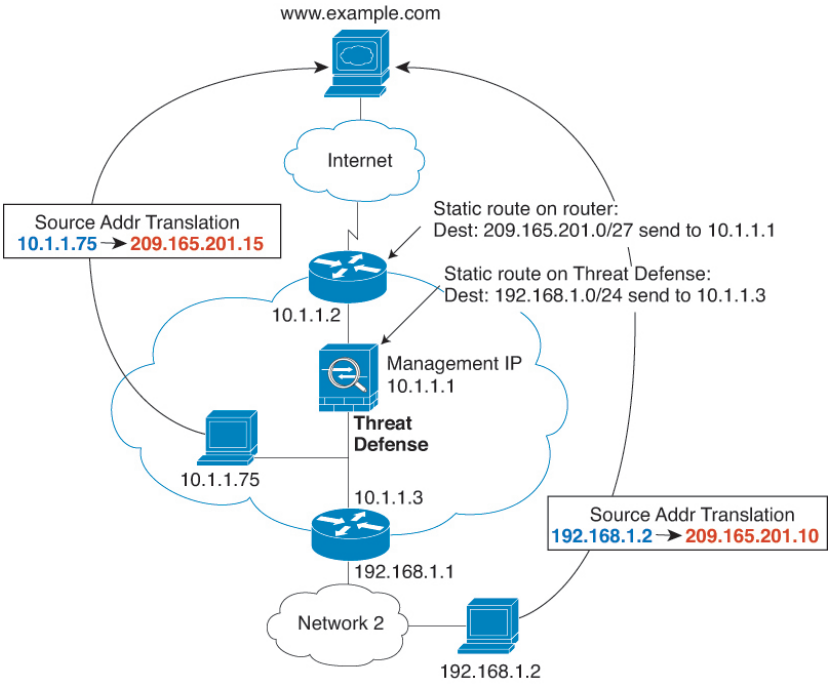
Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks. It can perform a similar function within a bridge group in routed mode.

NAT in transparent mode, or in routed mode between members of the same bridge group, has the following requirements and limitations:

- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the Firewall Threat Defense sends an ARP request to a host on the other side of the Firewall Threat Defense, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.
- Translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

The following figure shows a typical NAT scenario in transparent mode, with the same network on the inside and outside interfaces. The transparent firewall in this scenario is performing the NAT service so that the upstream router does not have to perform NAT.

Figure 345: NAT Example: Transparent Mode



1. When the inside host at 10.1.1.75 sends a packet to a web server, the real source address of the packet, 10.1.1.75, is changed to a mapped address, 209.165.201.15.

2. When the server responds, it sends the response to the mapped address, 209.165.201.15, and the Firewall Threat Defense receives the packet because the upstream router includes this mapped network in a static route directed to the Firewall Threat Defense management IP address.
3. The Firewall Threat Defense then undoes the translation of the mapped address, 209.165.201.15, back to the real address, 10.1.1.1.75. Because the real address is directly-connected, the Firewall Threat Defense sends it directly to the host.
4. For host 192.168.1.2, the same process occurs, except for returning traffic, the Firewall Threat Defense looks up the route in its routing table and sends the packet to the downstream router at 10.1.1.3 based on the Firewall Threat Defense static route for 192.168.1.0/24.

Auto NAT and Manual NAT

You can implement address translation in two ways: *auto NAT* and *manual NAT*.

We recommend using auto NAT unless you need the extra features that manual NAT provides. It is easier to configure auto NAT, and it might be more reliable for applications such as Voice over IP (VoIP). (For VoIP, you might see a failure in the translation of indirect addresses that do not belong to either of the objects used in the rule.)

Auto NAT

All NAT rules that are configured as a parameter of a network object are considered to be *auto NAT* rules. This is a quick and easy way to configure NAT for a network object. You cannot create these rules for a group object, however.

Although these rules are configured as part of the object itself, you cannot see the NAT configuration in the object definition through the object manager.

When a packet enters an interface, both the source and destination IP addresses are checked against the auto NAT rules. The source and destination address in the packet can be translated by separate rules if separate matches are made. These rules are not tied to each other; different combinations of rules can be used depending on the traffic.

Because the rules are never paired, you cannot specify that sourceA/destinationA should have a different translation than sourceA/destinationB. Use manual NAT for that kind of functionality, where you can identify the source and destination address in a single rule.

Manual NAT

Manual NAT lets you identify both the source and destination address in a single rule. Specifying both the source and destination addresses lets you specify that sourceA/destinationA can have a different translation than sourceA/destinationB.

Note For static NAT, the rule is bidirectional, so be aware that “source” and “destination” are used in commands and descriptions throughout this guide even though a given connection might originate at the “destination” address. For example, if you configure static NAT with port address translation, and specify the source address as a Telnet server, and you want all traffic going to that Telnet server to have the port translated from 2323 to 23, then you must specify the *source* ports to be translated (real: 23, mapped: 2323). You specify the source ports because you specified the Telnet server address as the source address.

The destination address is optional. If you specify the destination address, you can either map it to itself (identity NAT), or you can map it to a different address. The destination mapping is always a static mapping.

Comparing Auto NAT and Manual NAT

The main differences between these two NAT types are:

- How you define the real address.
  - Auto NAT—The NAT rule becomes a parameter for a network object. The network object IP address serves as the original (real) address.
  - Manual NAT—You identify a network object or network object group for both the real and mapped addresses. In this case, NAT is not a parameter of the network object; the network object or group is a parameter of the NAT configuration. The ability to use a network object *group* for the real address means that manual NAT is more scalable.
- How source and destination NAT is implemented.
  - Auto NAT— Each rule can apply to either the source or destination of a packet. So two rules might be used, one for the source IP address, and one for the destination IP address. These two rules cannot be tied together to enforce a specific translation for a source/destination combination.
  - Manual NAT—A single rule translates both the source and destination. A packet matches one rule only, and further rules are not checked. Even if you do not configure the optional destination address, a matching packet still matches one manual NAT rule only. The source and destination are tied together, so you can enforce different translations depending on the source/destination combination. For example, sourceA/destinationA can have a different translation than sourceA/destinationB.
- Order of NAT Rules.
  - Auto NAT—Automatically ordered in the NAT table.
  - Manual NAT—Manually ordered in the NAT table (before or after auto NAT rules).

NAT rule order

Auto NAT and manual NAT rules are stored in a single table that is divided into three sections. Section 1 rules are applied first, then section 2, and finally section 3, until a match is found. For example, if a match is found in section 1, sections 2 and 3 are not evaluated. The following table shows the order of rules within each section.



Note There is also a Section 0, which contains any NAT rules that the system creates for its own use. These rules have priority over all others. The system automatically creates these rules and clears xlates as needed. You cannot add, edit, or modify rules in Section 0.

Table 96: NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Manual NAT	Applied on a first match basis, in the order they appear in the configuration. Because the first match is applied, you must ensure that specific rules come before more general rules, or the specific rules might not be applied as desired. By default, manual NAT rules are added to section 1. By "specific rules first," we mean: • Static rules should come before dynamic rules. • Rules that include destination translation should come before rules with source translation only. If you cannot eliminate overlapping rules, where more than one rule might apply based on the source or destination address, be especially careful to follow these recommendations. If you do not follow the above recommendation, ensure that accurate routes are configured for the mapped addresses to match the correct NAT rule.
Section 2	Auto NAT	If a match in section 1 is not found, section 2 rules are applied in the following order: **1.** Static rules. **2.** Dynamic rules. **Within each rule type, the following ordering guidelines are used:** **1.** Quantity of real IP addresses—From smallest to largest. For example, an object with one address will be assessed before an object with 10 addresses. **2.** For quantities that are the same, then the IP address number is used, from lowest to highest. For example, 10.1.1.0 is assessed before 11.1.1.0. **3.** If the same IP address is used, then the name of the network object is used, in alphabetical order. For example, abracadabra is assessed before catwoman.
Section 3	Manual NAT	If a match is still not found, section 3 rules are applied on a first match basis, in the order they appear in the configuration. This section should contain your most general rules. You must also ensure that any specific rules in this section come before general rules that would otherwise apply.

For section 2 rules, for example, you have the following IP addresses defined within network objects:

- 192.168.1.0/24 (static)

- 192.168.1.0/24 (dynamic)
- 10.1.1.0/24 (static)
- 192.168.1.1/32 (static)
- 172.16.1.0/24 (dynamic) (object def)
- 172.16.1.0/24 (dynamic) (object abc)

The resultant ordering would be:

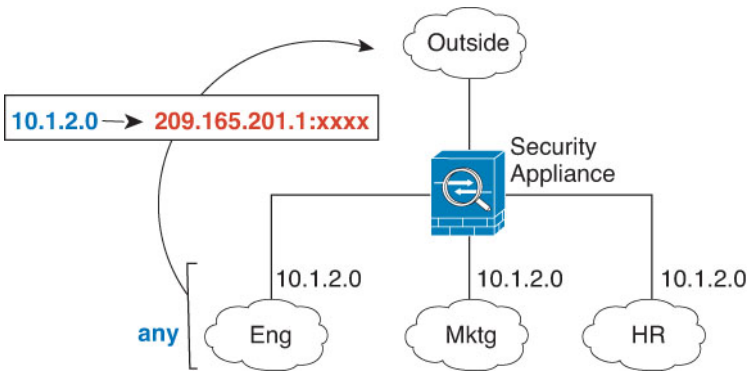
- 192.168.1.1/32 (static)
- 10.1.1.0/24 (static)
- 192.168.1.0/24 (static)
- 172.16.1.0/24 (dynamic) (object abc)
- 172.16.1.0/24 (dynamic) (object def)
- 192.168.1.0/24 (dynamic)

NAT interfaces

Except for bridge group member interfaces, you can configure a NAT rule to apply to any interface (in other words, all interfaces), or you can identify specific real and mapped interfaces. You can also specify any interface for the real address, and a specific interface for the mapped address, or vice versa.

For example, you might want to specify any interface for the real address and specify the outside interface for the mapped address if you use the same private addresses on multiple interfaces, and you want to translate them all to the same global pool when accessing the outside.

Figure 346: Specifying Any Interface



However, the concept of "any" interface does not apply to bridge group member interfaces. When you specify "any" interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. This could result in many similar rules where only one interface is different. You cannot configure NAT for the Bridge Virtual Interface (BVI) itself, you can configure NAT for member interfaces only.

Note You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes. When specifying interfaces, you do so indirectly by selecting the interface object that contains the interface.

NAT exemption

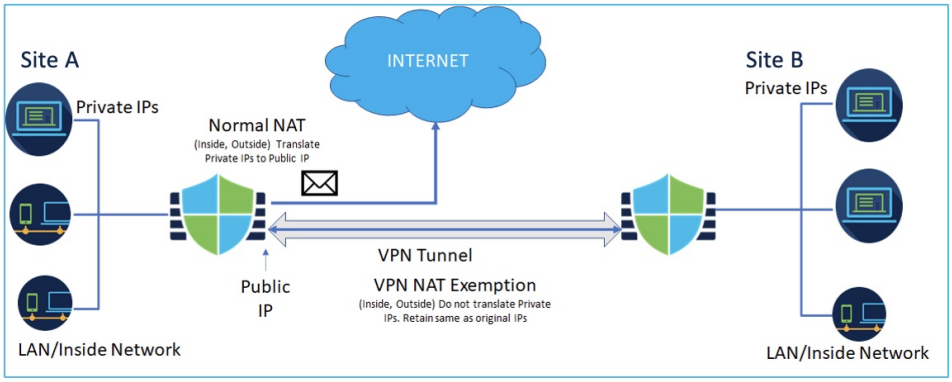
When internet edge devices have a site-to-site VPN configured on an interface and also have NAT rules for that interface, you must exempt the VPN traffic from the NAT rules. If you do not exempt the VPN traffic from NAT translation, the traffic gets dropped or is not routed through the VPN tunnel to the remote peer.

NAT exemption allows you to exclude traffic from being translated by NAT rules. When you create a policy-based site-to-site VPN using the Cloud-Delivered Firewall Management Center VPN wizard, you can select the **NAT Exempt** option to create the rules automatically (**Devices** > **VPN** > **Site to Site**). By default, this option is enabled. You can view the NAT exemptions for a device in the NAT policy page (**Devices** > **NAT**, and then click **NAT Exemptions**).

The Cloud-Delivered Firewall Management Center supports NAT exemption for all policy-based site-to-site VPN topology types. For more information, see Configure a policy-based site-to-site VPN, on page 1731.

Consider the following example, which shows a site-to-site VPN tunnel connecting Site A and Site B. For traffic that must go to the Internet, NAT translates the private IPs to a public IP address to access the Internet. For traffic that must go over the VPN tunnel, you must configure NAT exemption for the device in the VPN wizard.

Figure 347: Site-to-site VPN Topology with NAT Exemption



Note that when you upgrade from versions earlier than 7.4, if you have enabled NAT exemption, Cloud-Delivered Firewall Management Center will disable this option. You must enable this option in the policy-based site-to-site VPN wizard.

Configuring routing for NAT

The Firewall Threat Defense device needs to be the destination for any packets sent to the translated (mapped) address.

When sending packets, the device uses the destination interface if you specify one, or a routing table lookup if you do not, to determine the egress interface. For identity NAT, you have the option to use a route lookup even if you specify a destination interface.

The type of routing configuration needed depends on the type of mapped address, as explained in the following topics.

Addresses on the same network as the mapped interface

If you use addresses on the same network as the destination (mapped) interface, the Firewall Threat Defense uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the Firewall Threat Defense does not have to be the gateway for any additional networks. This solution is ideal if the outside network contains an adequate number of free addresses, a consideration if you are using a 1:1 translation like dynamic NAT or static NAT. Dynamic PAT greatly extends the number of translations you can use with a small number of addresses, so even if the available addresses on the outside network is small, this method can be used. For PAT, you can even use the IP address of the mapped interface.



Note If you configure the mapped interface to be any interface, and you specify a mapped address on the same network as one of the mapped interfaces, then if an ARP request for that mapped address comes in on a *different* interface, then you need to manually configure an ARP entry for that network on the ingress interface, specifying its MAC address. Typically, if you specify any interface for the mapped interface, then you use a unique network for the mapped addresses, so this situation would not occur. Configure the ARP table in the ingress interface's **Advanced** settings.

Addresses on a unique network

If you need more addresses than are available on the destination (mapped) interface network, you can identify addresses on a different subnet. The upstream router needs a static route for the mapped addresses that points to the Firewall Threat Defense.

Alternatively for routed mode, you can configure a static route on the Firewall Threat Defense for the mapped addresses using any IP address on the destination network as the gateway, and then redistribute the route using your routing protocol. For example, if you use NAT for the inside network (10.1.1.0/24) and use the mapped IP address 209.165.201.5, then you can configure a static route for 209.165.201.5 255.255.255.255 (host address) to the 10.1.1.99 gateway that can be redistributed.

For transparent mode, if the real host is directly-connected, configure the static route on the upstream router to point to the Firewall Threat Defense: specify the bridge group IP address. For remote hosts in transparent mode, in the static route on the upstream router, you can alternatively specify the downstream router IP address.

The same address as the real address (identity NAT)

The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. You can also disable proxy ARP for regular static NAT if desired, in which case you need to be sure to have proper routes on the upstream router.

Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues. For example, if you configure a broad identity NAT rule for “any” IP address, then leaving proxy ARP enabled can cause problems for hosts on the network directly connected to the mapped interface. In this case, when a host on the mapped network wants to communicate with another host on the same network, then the address in the ARP request matches the NAT rule (which matches “any” address). The Firewall Threat Defense will then proxy ARP for the address, even though the packet is not actually destined for the Firewall Threat Defense. (Note that this problem occurs even if you have a manual NAT rule; although the NAT rule must match both the source and destination addresses, the proxy ARP decision is made only on the “source” address). If the

Firewall Threat Defense ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the Firewall Threat Defense.

Requirements and prerequisites for NAT policies

Ensure you meet these requirements for configuring NAT policies:

User roles

You must have one of these user privileges to configure and use NAT policies:

- Admin
- Access Admin
- Network Admin

Guidelines for NAT

The following topics provide detailed guidelines for implementing NAT.

Firewall mode guidelines for NAT

NAT is supported in routed and transparent firewall mode.

However, configuring NAT on bridge group member interfaces (interfaces that are part of a Bridge Group Virtual Interface, or BVI) has the following restrictions:

- When configuring NAT for the members of a bridge group, you specify the member interface. You cannot configure NAT for the bridge group interface (BVI) itself.
- When doing NAT between bridge group member interfaces, you must specify the real and mapped addresses. You cannot specify “any” as the interface.
- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- You cannot translate between IPv4 and IPv6 networks (NAT64/46) when the source and destination interfaces are members of the same bridge group. Static NAT/PAT 44/66, dynamic NAT44/66, and dynamic PAT44 are the only allowed methods; dynamic PAT66 is not supported. However, you can do NAT64/46 between members of different bridge groups, or between a bridge group member (source) and standard routed interface (destination).

Note You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes.

IPv6 NAT guidelines

NAT supports IPv6 with the following guidelines and restrictions.

- For standard routed mode interfaces, you can also translate between IPv4 and IPv6.
- You cannot translate between IPv4 and IPv6 for interfaces that are members of the same bridge group. You can translate between two IPv6 or two IPv4 networks only. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- For static NAT, you can specify an IPv6 subnet up to /64. Larger subnets are not supported.
- When using FTP with NAT46, when an IPv4 FTP client connects to an IPv6 FTP server, the client must use either the extended passive mode (EPSV) or extended port mode (EPRT); PASV and PORT commands are not supported with IPv6.

IPv6 NAT best practices

You can use NAT to translate between IPv6 networks, and also to translate between IPv4 and IPv6 networks (routed mode only). We recommend the following best practices:

- NAT66 (IPv6-to-IPv6)—We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (manual NAT only).
- NAT46 (IPv4-to-IPv6)—We recommend using static NAT. Because the IPv6 address space is so much larger than the IPv4 address space, you can easily accommodate a static translation. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (manual NAT only). When translating to an IPv6 subnet (/96 or lower), the resulting mapped address is by default an IPv4-embedded IPv6 address, where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For example, if the IPv6 prefix is a /96 prefix, then the IPv4 address is appended in the last 32-bits of the address. For example, if you map 192.168.1.0/24 to 201b::0/96, then 192.168.1.4 will be mapped to 201b::0.192.168.1.4 (shown with mixed notation). If the prefix is smaller, such as /64, then the IPv4 address is appended after the prefix, and a suffix of 0s is appended after the IPv4 address. You can also optionally translate the addresses net-to-net, where the first IPv4 address maps to the first IPv6 address, the second to the second, and so on.
- NAT64 (IPv6-to-IPv4)—You may not have enough IPv4 addresses to accommodate the number of IPv6 addresses. We recommend using a dynamic PAT pool to provide a large number of IPv4 translations.

NAT support for inspected protocols

Some application layer protocols that open secondary connections, or that embedded IP addresses in packets, are inspected to provide the following services:

- Pinhole creation—Some application protocols open secondary TCP or UDP connections either on standard or negotiated ports. Inspection opens pinholes for these secondary ports so that you do not need to create access control rules to allow them.

- NAT rewrite— Protocols such as FTP embed IP addresses and ports for the secondary connections in packet data as part of the protocol. If there is NAT translation involved for either of the endpoints, the inspection engines rewrite the packet data to reflect the NAT translation of the embedded addresses and ports. The secondary connections would not work without NAT rewrite.
- Protocol enforcement—Some inspections enforce some degree of conformance to the RFCs for the inspected protocol.

The following table lists the inspected protocols that apply NAT rewrite and their NAT limitations. Keep these limitations in mind when writing NAT rules that include these protocols. Inspected protocols not listed here do not apply NAT rewrite. These inspections include GTP, HTTP, IMAP, POP, SMTP, SSH, and SSL.

Note NAT rewrite is supported on the listed ports only. For some of these protocols, you can extend inspection to other ports using Network Analysis Policies, but NAT rewrite is not extended to those ports. This includes DCERPC, DNS, FTP, and Sun RPC inspection. If you use these protocols on non-standard ports, do not use NAT on the connections.

Table 97: NAT Supported Application Inspection

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
DCERPC	TCP/135	No NAT64.	Yes
DNS over UDP	UDP/53	No NAT support is available for name resolution through WINS.	No
ESMTP	TCP/25	No NAT64.	No
FTP	TCP/21	(Clustering) No static PAT.	Yes
H.323 H.225 (Call signaling) H.323 RAS	TCP/1720 UDP/1718 For RAS, UDP/1718-1719	(Clustering) No static PAT. No extended PAT. No NAT64.	Yes
ICMP ICMP Error	ICMP (ICMP traffic directed to a device interface is never inspected.)	No limitations.	No
IP Options	RSVP	No NAT64.	No
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	No extended PAT. No NAT64.	No
RSH	TCP/514	No PAT. No NAT64. (Clustering) No static PAT.	Yes

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
RTSP	TCP/554 (No handling for HTTP cloaking.)	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes
SIP	TCP/5060 UDP/5060	No extended PAT. No NAT64 or NAT46. (Clustering) No static PAT.	Yes
Skinny (SCCP)	TCP/2000	No extended PAT. No NAT64, NAT46, or NAT66. (Clustering) No static PAT.	Yes
SQL*Net (versions 1, 2)	TCP/1521	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes
Sun RPC	TCP/111 UDP/111	No extended PAT. No NAT64.	Yes
TFTP	UDP/69	No NAT64. (Clustering) No static PAT. Payload IP addresses are not translated.	Yes
XDMCP	UDP/177	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes

FQDN destination guidelines

You can specify the translated (mapped) destination in a manual NAT rule using a fully-qualified domain name (FQDN) network object instead of an IP address. For example, you can create a rule based on traffic that is destined for the www.example.com web server.

When using an FQDN, the system obtains the DNS resolution and writes the NAT rule based on the returned address. If you are using multiple DNS server groups, the filter domains are honored and the address is requested from the appropriate group based on the filters. If more than one address is obtained from the DNS server, the address used is based on the following:

- If there is an address on the same subnet as the specified interface, that address is used. If there isn't one on the same subnet, the first address returned is used.
- The IP type for the translated source and translated destination must match. For example, if the translated source address is IPv6, the FQDN object must specify IPv6 as the address type. If the translated source is IPv4, the FQDN object can specify IPv4 or both IPv4 and IPv6. In this case, an IPv4 address is selected.

You cannot include an FQDN object in a network group that is used for manual NAT destination. In NAT, an FQDN object must be used alone, as only a single destination host makes sense for this type of NAT rule.

If the FQDN cannot be resolved to an IP address, the rule is not functional until a DNS resolution is obtained.

Additional guidelines for NAT

- NAT rules apply to through the device traffic only, they do not apply to traffic initiated by the device, such as a RADIUS authentication.
- For interfaces that are members of a bridge group, you write NAT rules for the member interfaces. You cannot write NAT rules for the Bridge Virtual Interface (BVI) itself.
- You cannot write NAT rules for a Virtual Tunnel Interface (VTI), which are used in site-to-site VPN. Writing rules for the VTI's source interface will not apply NAT to the VPN tunnel. To write NAT rules that will apply to VPN traffic tunneled on a VTI, you must use "any" as the interface; you cannot explicitly specify interface names.
- (Auto NAT only.) You can only define a single NAT rule for a given object; if you want to configure multiple NAT rules for an object, you need to create multiple objects with different names that specify the same IP address.
- If a VPN is defined on an interface, inbound ESP traffic on the interface is not subject to the NAT rules. The system allows the ESP traffic for established VPN tunnels only, dropping traffic not associated with an existing tunnel. This restriction applies to ESP and UDP ports 500 and 4500.
- If you define a site-to-site VPN on a device that is behind a device that is applying dynamic PAT, so that UDP ports 500 and 4500 are not the ones actually used, you must initiate the connection from the device that is behind the PAT device. The responder cannot initiate the security association (SA) because it does not know the correct port numbers.
- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT configuration is used, you can clear the translation table using the **clear xlate** command in the device CLI. However, clearing the translation table disconnects all current connections that use translations.

  If you create a new NAT rule that should apply to an existing connection (such as a VPN tunnel), you need to use **clear conn** to end the connection. Then, the attempt to re-establish the connection should hit the NAT rule and the connection should be NAT'ed correctly.



Note If you remove a dynamic NAT or PAT rule, and then add a new rule with mapped addresses that overlap the addresses in the removed rule, then the new rule will not be used until all connections associated with the removed rule time out or are cleared using the **clear xlate** or **clear conn** commands. This safeguard ensures that the same address is not assigned to multiple hosts.

- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.
- A network object used in NAT cannot include more than 131,838 IP addresses, either explicitly or implied in a range of addresses or a subnet. Break up the address space into smaller ranges and write separate rules for the smaller objects.

- (Manual NAT only.) When using **any** as the source address in a NAT rule, the definition of “any” traffic (IPv4 vs. IPv6) depends on the rule. Before the Firewall Threat Defense performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the Firewall Threat Defense can determine the value of **any** in a NAT rule. For example, if you configure a rule from “any” to an IPv6 server, and that server was mapped from an IPv4 address, then **any** means “any IPv6 traffic.” If you configure a rule from “any” to “any,” and you map the source to the interface IPv4 address, then **any** means “any IPv4 traffic” because the mapped interface address implies that the destination is also IPv4.
- You can use the same mapped object or group in multiple NAT rules.
- The mapped IP address pool cannot include:
  - The mapped interface IP address. If you specify “any” interface for the rule, then all interface IP addresses are disallowed. For interface PAT (routed mode only), specify the interface name instead of the interface address.
  - The failover interface IP address.
  - (Transparent mode.) The management IP address.
  - (Dynamic NAT.) The standby interface IP address when VPN is enabled.
- Avoid using overlapping addresses in static and dynamic NAT policies. For example, with overlapping addresses, a PPTP connection can fail to get established if the secondary connection for PPTP hits the static instead of dynamic xlate.
- You cannot use overlapping addresses in the source address of a NAT rule and a remote access VPN address pool.
- If you specify a destination interface in a rule, then that interface is used as the egress interface rather than looking up the route in the routing table. However, for identity NAT, you have the option to use a route lookup instead.
- If you use PAT on Sun RPC traffic, which is used to connect to NFS servers, be aware that the NFS server might reject connections if the PAT'ed port is above 1024. The default configuration of NFS servers is to reject connections from ports higher than 1024. The error is typically "Permission Denied." Mapping ports above 1024 happens if you do not select the option to include the reserved ports (1-1023) in the port range of a PAT pool. You can avoid this problem by changing the NFS server configuration to allow all port numbers.
- NAT applies to through traffic only. Traffic generated by the system is not subject to NAT.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.
- You cannot use NAT on the internal payload of Protocol Independent Multicast (PIM) registers.
- (Manual NAT) When writing NAT rules for a dual ISP interface setup (primary and backup interfaces using service level agreements in the routing configuration), do not specify destination criteria in the rule. Ensure the rule for the primary interface comes before the rule for the backup interface. This allows the device to choose the correct NAT destination interface based on the current routing state when the primary ISP is unavailable. If you specify destination objects, the NAT rule will always select the primary interface for the otherwise duplicate rules.

- If you get the ASP drop reason nat-no-xlate-to-pat-pool for traffic that should not match the NAT rules defined for the interface, configure identity NAT rules for the affected traffic so the traffic can pass untranslated.
- If you configure NAT for GRE tunnel endpoints, you must disable keepalives on the endpoints or the tunnel cannot be established. The endpoints send keepalives to the original addresses.
- DHCP and BOOTP share ports UDP/67-68. Because BOOTP is obsolete, writing NAT rules for the bootps port can cause port allocation problems when also running DHCP. Consider using DHCP relay instead for transmitting DHCP requests between network segments.
- (Secure Firewall 200) If you use PAT with SIP sessions, and you send thousands of SIP connections through the device, reduce the SIP invite timeout to the minimum of 1 minute. This can avoid problems with free memory consumption by xlates related to the SIP sessions. Configure the timeout on the **Timeout** page in the Platform Settings policy.
- In rare cases, return traffic (server to client) with an existing translation (xlate) might be logged as a new flow in Connection Events. This can occur when the client has already terminated the connection and the server sends another packet that reaches the device during the brief interval between connection closure and xlate removal, often due to application behavior or TCP stack cleanup. Because the device removes the xlate only after deleting the connection, a server packet can arrive while the xlate still exists. If no valid connection entry is found, the device logs a separate connection event based on the matched Access Control Policy rule.

Manage NAT policies

NAT converts the IP address of an incoming packet to a different address in the outgoing packet. One of the main functions of NAT is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into routable addresses that can be used on the public Internet. NAT keeps track of the translations, also known as xlates, to ensure that return traffic is directed to the correct untranslated host address.

Procedure

Step 1 Choose **Devices** > **NAT**.

Step 2 Manage your NAT policies:

- Create—Click **New Policy** and select **Threat Defense NAT**. Refer to Create NAT policies, on page 1584.
- Copy—Click **Copy** ( ) next to the policy you want to copy. You are prompted to give the copy a new, unique name. The copy includes all policy rules and configurations, but does not include device assignments.
- Report—Click **Report** ( ) for the policy. You are prompted to save the PDF report, which includes policy attributes, device assignments, rules, and object usage information.
- Edit—Click **Edit** ( ) next to the policy you want to edit. Refer to Configure NAT for threat defense, on page 1585.

- Delete—Click **Delete** ( ) next to the policy you want to delete, then click **OK**. When prompted whether to continue, you are also informed if another user has unsaved changes in the policy.

  Caution
  After you have deployed a NAT policy to a managed device, you cannot delete the policy from the device. Instead, you must deploy a NAT policy with no rules to remove the NAT rules already present on the managed device. You also cannot delete a policy that is the last deployed policy on any of its target devices, even if it is out of date. Before you can delete the policy completely, you must deploy a different policy to those targets.

Create NAT policies

Create a NAT policy to define how network addresses are translated between different network segments.

When you create a new NAT policy you must, at minimum, give it a unique name. Although you are not required to identify policy targets at policy creation time, you must perform this step before you can deploy the policy. If you apply a NAT policy with no rules to a device, the system removes all NAT rules from that device.

Procedure

Step 1 Choose **Devices** > **NAT**.

Step 2 Click **New Policy** and from the drop-down list, choose **Threat Defense NAT** for Firewall Threat Defense devices.

Firepower NAT is for older devices that are not covered in this document.

Step 3 Enter a unique **Name**.

Step 4 (Optional) Enter a **Description**.

Step 5 Choose the devices where you want to deploy the policy:

- Choose a device in the **Available Devices** list, and then click **Add to Policy**.

  Alternatively, click and drag a device from the **Available Devices** list to the **Selected Devices** list.

- Remove a device from the **Selected Devices** list by clicking **Delete** ( ) next to the device.

Step 6 Click **Save**.

Configure NAT policy targets

Identify the managed devices that you want to target with your policy while creating or editing a policy and assign NAT policies to identified devices to control NAT behavior. You can search a list of available devices and high-availability pairs, and add them to a list of selected devices.

Procedure

Step 1 Choose **Devices** > **NAT**.

Step 2 Click **Edit** ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Policy Assignments**.

Step 4 Do any of the following:

- To assign a device, high-availability pair, or device group to the policy, select it in the **Available Devices** list and click **Add to Policy**. Alternatively, you can also drag and drop selected devices.
- To remove a device assignment, click **Delete** ( ) next to a device, high-availability pair, or device group in the **Selected Devices** list.

Step 5 Click **OK**.

Configure NAT for threat defense

The NAT policy is a shared policy. Assign it to devices that need similar NAT rules.

Each rule applies to a device if its interface objects (security zones or interface groups) include an interface for that device. When the interface objects include one or more interface for the device, the rule is deployed to the device. You can target subsets of devices within a shared policy by designing interface objects accordingly. Rules using the "any" interface object apply to all devices.

If you change an interface type to one not valid for NAT on an assigned device, the policy marks the interface as deleted. Click **Save** in the NAT policy to automatically remove the interface from the policy.

You can configure multiple NAT policies if groups of your devices require significantly different rules.

Procedure

Step 1 Navigate from Security Cloud Control to Cloud-Delivered Firewall Management Center, on page 3094

Step 2 Select **Devices** > **NAT**.

- Click **New Policy** > **Threat Defense NAT** to create a new policy. Give the policy a name, optionally assign devices to it, and click **Save**.

  You can change device assignments later by editing the policy and clicking **Policy Assignments**.

- Click **Edit** ( ) to edit an existing threat defense NAT policy. Note that the page also shows Firepower NAT policies, which are not used by Firewall Threat Defense devices.

  If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Decide what kinds of rules you need.

You can create dynamic NAT, dynamic PAT, static NAT, and identity NAT rules. For an overview, refer to NAT types, on page 1568.

Step 4 Decide which rules should be implemented as manual or auto NAT.

For a comparison of these two implementation options, refer to Auto NAT and Manual NAT, on page 1571.

Step 5 Decide which rules should be custom per device.

Because you can assign a NAT policy to multiple devices, you can configure a single rule on many devices. However, you might have rules that should be interpreted differently by each device, or some rules that should apply to a subset of devices only.

Use interface objects to control on which devices a rule is configured. Then, use object overrides on network objects to customize the addresses used per device.

For detailed information, refer to Customize NAT rules for multiple devices, on page 1587.

Step 6 Create the rules as explained in these sections.

- Dynamic NAT, on page 1593
- Dynamic PAT, on page 1598
- Static NAT, on page 1608
- Identity NAT, on page 1617

Step 7 Do these to manage your NAT policy:

- To edit the policy name or description, click in those fields, type in your changes, and click outside the fields.
- To view only those rules that apply to a specific device, click **Filter by Device** and select the desired device. A rule applies to a device if it uses an interface object that includes an interface on the device.
- To view any warnings or errors in the policy, click **Show Warnings**, then choose a **Device**. Warnings and errors mark configurations that could adversely affect traffic flow or prevent the policy from deploying.
- To change the devices to which the policy is assigned, click the **Policy Assignments** link and modify the selected devices list as desired.
- To change the number of rules displayed on the page, use the **Rows Per Page** drop-down list.

Step 8 Do these to manage your NAT rules:

- To add a rule, click the **Add Rule** button.
- To edit a rule, click the **Edit** ( ) icon next to the rule.
- To delete a rule, click the **Delete** ( ) icon next to the rule.
- To enable or disable a rule, right-click the rule and choose **Enable** or **Disable** from the **State** menu. You can temporarily disable a rule without deleting it using these controls.
- To select more than one rule to enable, disable, or delete, click the checkbox for the rules, or the checkbox in the header, then perform the action.

Step 9 Click **Save**.

You can now go to **Deploy** > **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Customize NAT rules for multiple devices

Because the NAT policy is shared, you can assign a given policy to more than one device. However, you can configure at most one auto NAT rule for a given object. Thus, if you want to configure different translations for an object based on the specific device doing the translation, you need to carefully configure the interface objects (security zones or interface groups) and define network object overrides for the translated address. The interface objects determine on which devices a rule gets configured. The network object overrides determine what IP addresses are used by a given device for that object.

Consider the following scenario:

- FTD-A and FTD-B have inside networks 192.168.1.0/24 attached to the interface named “inside.”
- On FTD-A, you want to translate all 192.168.1.0/24 addresses to a NAT pool in the 10.100.10.10 - 10.100.10.200 range when going to the “outside” interface.
- On FTD-B, you want to translate all 192.168.1.0/24 addresses to a NAT pool in the 10.200.10.10 - 10.200.10.200 range when going to the “outside” interface.

Do these steps to accomplish the above scenario. Although this example rule is for dynamic auto NAT, you can generalize the technique for any type of NAT rule.

Procedure

Step 1 Create the security zones for the inside and outside interfaces.

a) Choose **Objects** > **Object Management**.
b) Select **Interface** from the table of contents and click **Add** > **Security Zone**.

You can use interface groups instead of zones.

c) Configure the inside zone properties.

- **Name**—Enter a name, for example, **inside-zone**.
- **Type**—Select **Routed** for routed-mode devices, **Switched** for transparent mode.
- **Selected Interfaces**—Add the FTD-A/inside and FTD-B/inside interfaces to the selected list.

d) Click **Save**.
e) Click **Add** > **Security Zone** and define the outside zone properties.

- **Name**—Enter a name, for example, **outside-zone**.
- **Interface Type**—Select **Routed** for routed-mode devices, **Switched** for transparent mode.
- **Selected Interfaces**—Add the FTD-A/outside and FTD-B/outside interfaces to the selected list.

f) Click **Save**.

Step 2 Create the network object for the original inside network on the Object Management page.

a) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

b) Configure the inside network properties.

- **Name**—Enter a name, for example, **inside-network**.
- **Network**—Enter the network address, for example, **192.168.1.0/24**.

c) Click **Save**.

Step 3 Create the network object for the translated NAT pool and define overrides.

a) Click **Add Network** > **Add Object**.

b) Configure the NAT pool properties for FTD-A.

- **Name**—Enter a name, for example, **NAT-pool**.
- **Network**—Enter the range of addresses to include in the pool for FTD-A, for example, **10.100.10.10-10.100.10.200**.

c) Select **Allow Overrides**.

d) Click the **Overrides** heading to open the list of object overrides.

e) Click **Add** to open the Add Object Override dialog box.

f) Select FTD-B and **Add** it to the **Selected Devices** list.

g) Click **Override** and change **Network** to **10.200.10.10-10.200.10.200**.

h) Click **Add** to add the override to the device.

By defining an override for FTD-B, whenever the system configures this object on FTD-B, it will use the override value instead of the value defined in the original object.

i) Click **Save**.

Step 4 Configure the NAT rule.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside-zone.
- **Destination Interface Objects** = outside-zone.

Note
The interface objects control on which devices the rule is configured. Because in this example the zones contain interfaces for FTD-A and FTD-B only, even if the NAT policy were assigned to additional devices, the rule would be deployed to those 2 devices only.

e) On **Translation**, configure the following:

- **Original Source** = inside-network object.
- **Translated Source** > **Address**= NAT-pool object.

f) Click **Save**.

You now have a single rule that will be interpreted differently for FTD-A and FTD-B, providing unique translations for the inside networks protected by each firewall.

Search and filter the NAT rule table

You can search and filter the NAT rule table to help you find rules that you need to modify or view. When you filter the table, only matching rules are shown. Note that although the rule numbers change to be sequentially 1, 2, and so forth, filtering does not change the actual rule number or the rule's location in the table relative to hidden rules. Filtering simply changes what you can see to help you locate rules that interest you.

When editing the NAT policy, you can use the fields above the table to do these types of search or filter:

- **Filter by Device**—Click **Filter by Device**, then select the devices whose rules you want to see and click **OK**. Whether a rule applies to a device is determined by the rule's interface constraints. If you specify a security zone or interface group for either the source or destination interface, the rule applies to a device if at least one interface for the device is in the zone or group. If a NAT rule applies to any source and any destination interface, then it applies to all devices.

  If you also do a text or multiple-attribute search, the results are constrained to the selected devices.

  To remove this filter, click **Filter by Device** and deselect the devices, or select **All**, and click **OK**.

- **Simple Text Search**—In the **Filter** box, type a string and press Enter. The string is compared to all values in the rules. For example, if you enter "network-object-1," which is the name of a network object, you would get rules that use the object in source, destination, and PAT pool attributes.

  For network and port objects, the string is also compared to the contents of the objects used in the rule. For example, if a PAT pool object includes the range 10.100.10.3-10.100.10.100, searching on either 10.100.10.3 or 10.100.10.100 (or a partial 10.100.10) will include rules that use that PAT pool object. However, the match must be exact: searching on 10.100.10.5 will not match this PAT pool object, even though the IP address is within the object's IP address range.

  To remove the filter, click the **x** on the right side of the Filter box.

- **Multiple-Attribute Search**—If a simple text search gives you too many hits, you can configure multiple values for the search. Click in the **Filter** box to open the list of attributes, then select or enter strings for the attributes you intend to search and click the **Filter** button. These attributes are the same as the ones you would configure within a NAT rule. The attributes are and'ed, so filtered results include only those rules that match all attributes you configured.

  - For binary attributes, such as the rule state (enabled/disabled), whether a PAT pool is configured (enabled/disabled), the direction of the rule (uni/bi), or rule type (static/dynamic), simply check or uncheck the boxes as appropriate. Select both boxes if you do not care about the attribute value. If you deselect both boxes, no rules will match the filter.

- For string attributes, type a full or partial string relevant to that attribute. These will be object names, either for security zones/interface groups, network objects, or port objects. It can also be the network or port object contents, which are matched the same way they are for simple text searches.

To remove the filter, click the **x** on the right side of the Filter box, or click in the Filter box to open the drop-down list, and click the Clear button.

Enable, disable, or delete multiple rules

You can enable or disable manual NAT rules, or delete any NAT rule, one by one. You can also select multiple rules and apply changes to all of them at once. Because enable/disable applies to manual NAT only, if you select a mix of rule types, you can delete them only. Note that when you enable or disable rules, it does not matter if you select some rules that were already enabled or disabled. For example, enabling an already enabled rule simply leaves the rule enabled.

Procedure

Step 1 Select **Devices** > **NAT**, and edit a threat defense NAT policy.

Step 2 (Optional) Filter the NAT rules to locate the ones you want to change.

Filtering is especially useful if you have a large NAT policy. For example, you could search on disabled rules to find ones that need to be enabled.

Step 3 Select the rules you want to change.

- Click the checkbox in the left column of the rule to select (or deselect) individual rules.
- Click the checkbox in the table header to select all the rules on the currently displayed page.

Your selection is preserved as you go from page to page. However, in practice, it makes most sense to perform your actions on the rules selected on a page before going to the next page.

Step 4 Perform the desired action. When selecting multiple rules, you are asked to confirm the action.

Note that these actions are also available on the right-click menu.

- To enable all rules, click **Select Bulk Action** > **Enable**.
- To disable all rules, click **Select Bulk Action** > **Disable**.
- To delete all rules, click **Select Bulk Action** > **Delete**.

NAT templates

A NAT template is a predefined configuration tool that simplifies complex NAT setups by offering predefined recommended configurations. It streamlines the NAT setup process, making it accessible and efficient.

NAT template use cases

NAT templates are designed to address specific scenarios. The available templates cover these use cases:

- Allow internal hosts to access the Internet.
- Allow external hosts to access internal resources.

These templates simplify the configuration process by providing predefined NAT flows tailored to these typical requirements.

After you enter the required information for each template, a visual representation of the NAT rule appears at the bottom of the page. This feature provides a clear view of the NAT rule during addition or modification, enhancing the user interface experience.

Create a NAT rule for internal hosts

When internal hosts need to access resources across the Internet, you can use this template. This template creates a dynamic Auto NAT rule that translates the source to the destination interface IP.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Alternatively, you can create objects while defining the NAT rule.

Configure an **Original Source**. This must be a network object (not a group), and can be a host, range, or subnet.

Procedure

Step 1 Select **Policies** > **Network Policies** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Click the **Add** drop-down button to create a new rule.

Step 3 Under **Use templates**, choose the **Allow internal hosts to access internet** option.

Step 4 On the page that appears, select the Source interface, Destination Interface and Original source.

- Source interface is the interface group or zone from which the internal hosts are connected or closest.
- Destination interface is the exit interface group or zone, which is connected to or closest to Internet.
- Original Source is the internal network or IP addresses from which need to reach the Internet.

Note

- The **Create** button in the original source drop-down can be used to create a network object for the original source.
- The **Create** button in the Source interface or Destination interface drop-down can be used to create interface group or zone object to be used.

Step 5 Click **Add Rule** to save the rule.

Note

- The **Preset network configuration** drop-down shows the configurations preset for this template. You cannot change these configurations.

- The **Preview configuration** panel lets you view a visual representation of the rule during configuration.

Create a NAT rule for external hosts

When external hosts need to access internal resources, this template creates a static auto NAT rule that translates the source to the destination interface IP. Use this configuration when you need to allow external access to internal servers or services.

Before you begin

Select **Objects** and create the network objects or groups needed in the rule. Alternatively, create objects while you define a NAT rule.

Configure an **Original Source**. This must be a network object (not a group), and can be a host, range, or subnet.

Procedure

Step 1 Select **Policies** > **Network Policies** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Click the **Add** drop-down button to create a new rule.

Step 3 Under **Use templates** select the **Allow external hosts to access internal resources** option.

Step 4 On the page that appears, select the **Source interface**, **Destination interface**, and **Original source**.

- Source interface is the interface group or zone on which the internal server is connected or closest.
- Destination interface is the exit interface group or zone which is connected or closest to Internet.
- Original source is the internal network or IP addresses of the internal server.

Note

- The **Create** button in the original source drop-down can be used to create a network object for the original source.
- The **Create** button in the Source interface or Destination interface drop-down can be used to create interface group or zone.

Step 5 Select the **Port type** as TCP or UDP.

Step 6 Configure the original and translated ports.

- **Original source port** is the port on which the internal server is listening.
- **Translated source** port is the port that will be used to connect to the internal server.

Step 7 Click on **Add Rule** to save the rule.

Note

- The **Preset network configuration** drop-down shows the configurations preset for this template. You cannot change these configurations.

- The **Preview configuration** panel lets you view a visual representation of the rule during configuration.

Dynamic NAT

The following topics explain dynamic NAT and how to configure it.

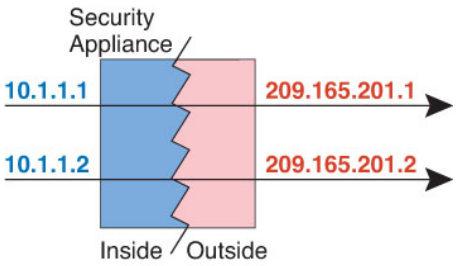
About dynamic NAT

Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool typically includes fewer addresses than the real group. When a host you want to translate accesses the destination network, NAT assigns the host an IP address from the mapped pool. The translation is created only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, even if the connection is allowed by an access rule.

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule. A successful connection from a remote host can reset the idle timer for the connection.

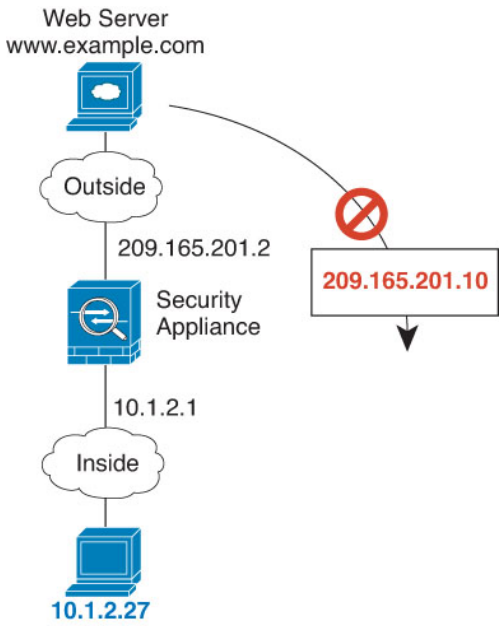
The following figure shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 348: Dynamic NAT



The following figure shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the packet is dropped.

Figure 349: Remote Host Attempts to Initiate a Connection to a Mapped Address



Dynamic NAT disadvantages and advantages

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

  Use PAT or a PAT fall-back method if this event occurs often because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool, and routable addresses may not be available in large quantities.

The advantage of dynamic NAT is that some protocols cannot use PAT. PAT does not work with the following:

- IP protocols that do not have a port to overload, such as GRE version 0.
- Some multimedia applications that have a data stream on one port, the control path on another port, and are not open standard.

Configure dynamic auto NAT

Use dynamic auto NAT rules to translate addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.

- **Translated Source**—This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Dynamic**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—The network object or group that contains the mapped addresses.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 1677.
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure dynamic manual NAT

Use dynamic manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic NAT translates addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—This can be a network object or group, but it cannot include a subnet. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule and select the **Manual NAT rule** option in the drop-down.
- Click **Edit** () to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

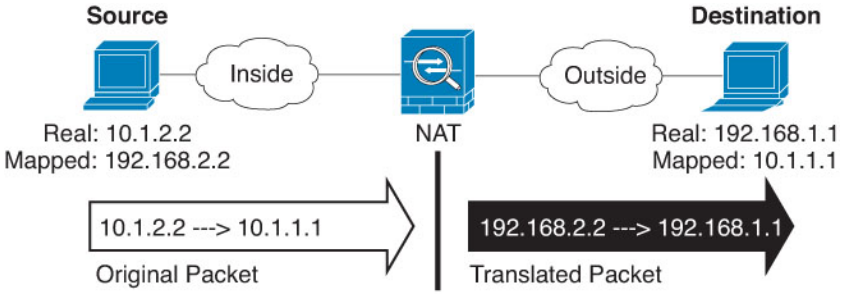
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—The network object or group that contains the mapped addresses.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

Step 8 (Optional.) On **Advanced**, select the desired options:

- (For source translation only.) **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 1677.
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

Dynamic PAT

The following topics describe dynamic PAT.

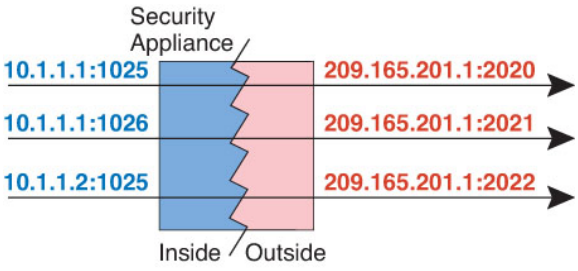
About dynamic PAT

Dynamic PAT translates multiple real addresses to a single mapped IP address by translating the real address and source port to the mapped address and a unique port.

Each connection requires a separate translation session because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

The following figure shows a typical dynamic PAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back. The mapped address is the same for each translation, but the port is dynamically assigned.

Figure 350: Dynamic PAT



For the duration of the translation, a remote host on the destination network can initiate a connection to the translated host if an access rule allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

After the connection expires, the port translation also expires.

Note We recommend that you use different PAT pools for each interface. If you use the same pool for multiple interfaces, especially if you use it for "any" interface, the pool can be quickly exhausted, with no ports available for new translations.

Dynamic PAT disadvantages and advantages

Dynamic PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the Firewall Threat Defense interface IP address as the PAT address.

You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. For more information, see NAT support for inspected protocols, on page 1578.

Dynamic PAT might also create a large number of connections appearing to come from a single IP address, and servers might interpret the traffic as a DoS attack. You can configure a PAT pool of addresses and use a round-robin assignment of PAT addresses to mitigate this situation.

PAT pool object guidelines

When creating network objects for a PAT pool, follow these guidelines.

For a PAT pool

- Ports are mapped to an available port in the 1024 to 65535 range. You can optionally include the reserved ports, those below 1024, to make the entire port range available for translations.

  When operating in a cluster, blocks of 512 ports per address are allocated to the members of the cluster, and mappings are made within these port blocks. If you also enable block allocation, the ports are distributed according to the block allocation size, whose default is also 512. If you change the cluster unit limit (the size of the cluster), ensure that you clear xlates or reboot the devices so that PAT pools can be appropriately reallocated to the cluster units.

- If you enable block allocation for a PAT pool, port blocks are allocated in the 1024-65535 range only. Thus, if an application requires a low port number (1-1023), it might not work. For example, an application requesting port 22 (SSH) will get a mapped port within the range of 1024-65535 and within the block allocated to the host.

- If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT, then the other rule must also specify extended PAT.

- If a host has an existing connection, then subsequent connections from that host use the same PAT IP address. If no ports are available, this can prevent the connection. Use the round robin option to avoid this problem.

- For best performance, limit the number of IP addresses within a PAT pool to 10,000.

For extended PAT for a PAT pool

- Many application inspections do not support extended PAT.
- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.
- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.
- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.
- You cannot use extended PAT on units in a cluster.
- Extended PAT increases memory usage on the device.

For round robin for a PAT pool

- If a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. However, this “stickiness” does not survive a failover. If the device fails over, then subsequent connections from a host might not use the initial IP address.
- IP address “stickiness” is also impacted if you mix PAT pool/round robin rules with interface PAT rules on the same interface. For any given interface, choose either a PAT pool or interface PAT; do not create competing PAT rules.
- Round robin, especially when combined with extended PAT, can consume a large amount of memory. Because NAT pools are created for every mapped protocol/IP address/port range, round robin results in a large number of concurrent NAT pools, which use memory. Extended PAT results in an even larger number of concurrent NAT pools.

Configure dynamic auto PAT

Use dynamic auto PAT rules to translate addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address (either the destination interface's address or another address), or use a PAT pool of addresses to provide a larger number of possible translations.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object.
  - **Single PAT address**—Create a network object containing a single host.
  - **PAT pool**—Create a network object that includes a range, or create a network object group that contains hosts, ranges, or both. You cannot include subnets. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** drop-down button to create a new rule and select the **Auto NAT rule** option in the drop-down.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Dynamic**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Skip the step for configuring a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Skip the step for configuring a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Do not select the destination interface IP or a single address as the **Translated Source**.

Step 6 If you are using a PAT pool, select the **PAT Pool** page and do the following:

a) Select **Enable PAT pool**.

b) Select the network object group that contains the addresses for the pool in the **PAT** > **Address** field.

You can alternatively select **Destination Interface IP**, which is another way to implement interface PAT.

c) (Optional) Select the following options as needed:

- **Use Round Robin Allocation**—To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.

- **Extended PAT Table**—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.
- **Flat Port Range**, **Include Reserved Ports**—To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option. For the Firewall Threat Defense devices running version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.
- **Block Allocation**—To enable port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation is compatible with round robin, but you cannot use it with the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Step 7 (Optional.) On **Advanced**, select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option. You cannot select this option if you already configured interface PAT as the translated address or PAT pool.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 8 Click **Save** to add the rule.

Step 9 Click **Save** on the NAT page to save your changes.

Configure dynamic manual PAT

Use dynamic manual PAT rules when auto PAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic PAT translates addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address (either the destination interface's address or another address), or use a PAT pool of addresses to provide a larger number of possible translations.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object.
  - **Single PAT address**—Create a network object containing a single host.
  - **PAT pool**—Create a network object that includes a range, or create a network object group that contains hosts, ranges, or both. You cannot include subnets.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule and select the **Manual NAT rule** option in the drop-down.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

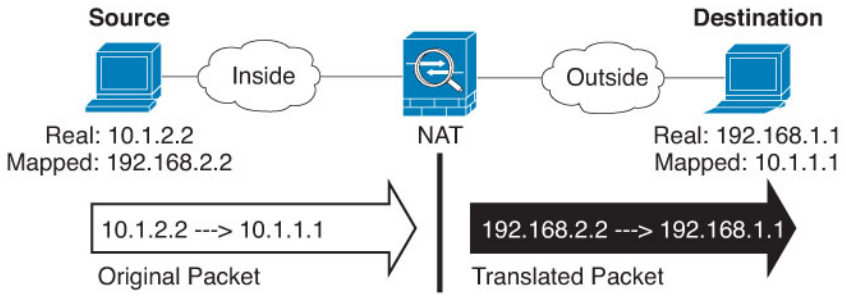
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Skip the step for configuring a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Skip the step for configuring a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Do not select the destination interface IP or a single address as the **Translated Source**.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

Step 8 If you are using a PAT pool, select the **PAT Pool** page and do the following:

a) Select **Enable PAT pool**.

b) Select the network object group that contains the addresses for the pool in the **PAT** > **Address** field.

You can alternatively select **Destination Interface IP**, which is another way to implement interface PAT.

c) (Optional) Select the following options as needed:

- **Use Round Robin Allocation**—To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extended PAT Table**—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.
- **Flat Port Range**, **Include Reserved Ports**—To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option. For the Firewall Threat Defense devices running version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.
- **Block Allocation**—To enable port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation is compatible with round robin, but you cannot use it with the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Step 9 (Optional.) On **Advanced**, select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 10 Click **Save** to add the rule.

Step 11 Click **Save** on the NAT page to save your changes.

Configure PAT with port block allocation

For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). If you allocate a block of ports, subsequent connections from the host use new randomly-selected ports within the block. If necessary, additional blocks are allocated

if the host has active connections for all ports in the original block. Blocks are freed when the last xlate that uses a port in the block is removed.

The main reason for allocating port blocks is reduced logging. The port block allocation is logged, connections are logged, but xlates created within the port block are not logged. On the other hand, this makes log analysis more difficult.

Port blocks are allocated in the 1024-65535 range only. There are separate blocks for TCP, UDP, and ICMP connections, and these blocks can overlap. Thus, if an application requires a low port number (1-1023), it might not work. For example, an application requesting port 22 (SSH) will get a mapped port within the range of 1024-65535 and within the block allocated to the host. You can create a separate NAT rule that does not use block allocation for applications that use low port numbers; for twice NAT, ensure the rule comes before the block allocation rule.

Before you begin

Usage notes for NAT rules:

- You can include the **Use Round Robin Allocation** option, but you cannot include the options for extending PAT uniqueness, using a flat range, including the reserved ports, or falling through to interface PAT. Other source/destination address and port information is also allowed.
- As with all NAT changes, if you replace an existing rule, you must clear xlates related to the replaced rule to have the new rule take effect. You can clear them explicitly or simply wait for them to time out. However, when operating in a cluster, you must reload the cluster so that the units can distribute blocks correctly.



Note If you are switching between a regular PAT and block allocation PAT rule, for object NAT, you must first delete the rule, then clear xlates. You can then create the new object NAT rule. Otherwise, you will see pat-port-block-state-mismatch drops in the **show asp drop** output.

- When you form a cluster while the device that will become the control node has an existing NAT policy that uses PAT pools, PAT port block distribution might not initialize correctly. The control node does not own or distribute the port blocks, and the **show nat pool cluster summary** command might show all port blocks as *RESERVED* on every cluster node.
  - Before forming a cluster, remove any existing NAT policy that uses PAT pools from the device that will become the control node. After cluster formation is complete, reconfigure and deploy the NAT policy.
  - If the cluster is already formed and all port blocks appear as *RESERVED*, delete the affected NAT policy and deploy the change. Then re-create and deploy the policy. Use the **show nat pool cluster summary** command to verify that the port blocks are distributed correctly across the cluster nodes.
- For a given PAT pool, you must specify (or not specify) block allocation for all rules that use the pool. You cannot allocate blocks in one rule and not in another. PAT pools that overlap also cannot mix block allocation settings. You also cannot overlap static NAT with port translation rules with the pool.

Procedure

Step 1 (Optional.) Configure global PAT port block allocation settings.

There are a few global settings that control port block allocation. If you want to change the defaults for these options, you must configure a FlexConfig object and add it to your FlexConfig policy.

a) Select **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object** and create a new object.

b) Configure the block allocation size, which is the number of ports in each block.

xlate block-allocation size *value*

The range is 32-4096. The default is 512. Use the "no" form to return to the default.

If you do not use the default, ensure that the size you choose divides evenly into 64,512 (the number of ports in the 1024-65535 range). Otherwise, there will be ports that cannot be used. For example, if you specify 100, there will be 12 unused ports.

Using the default block-allocation size, a single PAT IP can serve 126 unique, simultaneous users (number of ports 64,512 divided by 512).

c) Configure the maximum blocks that can be allocated per host.

xlate block-allocation maximum-per-host *number*

The limit is per protocol, so a limit of 4 means at most 4 UDP blocks, 4 TCP blocks, and 4 ICMP blocks per host. The range is 1-8, the default is 4. Use the "no" form to return to the default.

d) (Optional.) Enable interim syslog generation.

xlate block-allocation pba-interim-logging *seconds*

By default, the system generates syslog messages during port block creation and deletion. If you enable interim logging, the system generates the following message at the interval you specify. The messages report all active port blocks allocated at that time, including the protocol (ICMP, TCP, UDP) and source and destination interface and IP address, and the port block. You can specify an interval from 21600-604800 seconds (6 hours to 7 days).

%ASA-6-305017: Pba-interim-logging: Active *protocol* block of ports for translation from *real_interface*:*real_host_ip* to *mapped_interface*:*mapped_ip_address*/*start_port_num*-*end_port_num*

Example:

The following example sets the block allocation size to 64, the per-host maximum to 8, and enables interim logging every 6 hours.

```
xlate block-allocation size 64
xlate block-allocation maximum-per-host 8
xlate block-allocation pba-interim-logging 21600
```

e) Select the following options in the FlexConfig object:

- **Deployment** = **Everytime**
- **Type** = **Append**

f) Click **Save** to create the FlexConfig object.

g) Select **Devices** > **FlexConfig**, and create or edit the FlexConfig policy that is assigned to the devices that need to have these settings adjusted.
h) Select your object in the available objects list and click > to move it to the selected objects list.
i) Click **Save**.

You can click **Preview Config**, select one of the target devices, and verify that the xlate commands appear correctly.

Step 2 Add NAT rules that use PAT pool port block allocation.

a) Select **Devices** > **NAT** and add or edit the threat defense NAT policy.
b) Add or edit a NAT rule and configure at least the following options.

- **Type** = **Dynamic**
- In **Translation** > **Original Source**, select the object that defines the source address.
- On **PAT Pool**, configure the following options:
  - Select **Enable PAT Pool**
  - In **PAT** > **Address**, select a network object or group that defines the pat pool.
  - Select the **Block Allocation** option.

c) Save your changes to the rule and to the NAT policy.

Static NAT

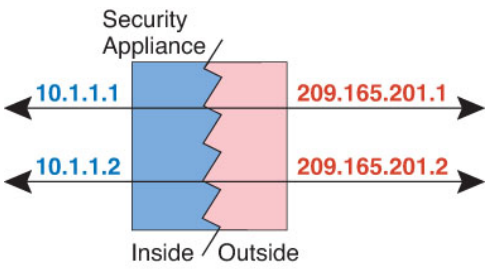
The following topics explain static NAT and how to implement it.

About static NAT

Static NAT creates a fixed translation of a real address to a mapped address. Because the mapped address is the same for each consecutive connection, static NAT allows bidirectional connection initiation, both to and from the host (if an access rule exists that allows it). With dynamic NAT and PAT, on the other hand, each host uses a different address or port for each subsequent translation, so bidirectional initiation is not supported.

The following figure shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 351: Static NAT



Note You can disable bidirectionality if desired.

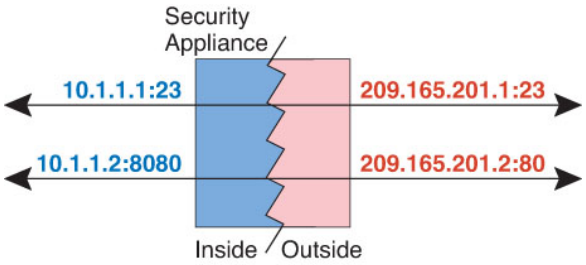
Static NAT with port translation

Static NAT with port translation lets you specify a real and mapped protocol and port.

When you specify the port with static NAT, you can choose to map the port and/or the IP address to the same value or to a different value.

The following figure shows a typical static NAT with port translation scenario showing both a port that is mapped to itself and a port that is mapped to a different value; the IP address is mapped to a different value in both cases. The translation is always active so both translated and remote hosts can initiate connections.

Figure 352: Typical Static NAT with Port Translation Scenario



Static NAT-with-port-translation rules limit access to the destination IP address for the specified port only. If you try to access the destination IP address on a different port not covered by a NAT rule, then the connection is blocked. In addition, for manual NAT, traffic that does not match the source IP address of the NAT rule will be dropped if it matches the destination IP address, regardless of the destination port. Therefore, you must add additional rules for all other traffic allowed to the destination IP address. For example, you can configure a static NAT rule for the IP address, without port specification, and place it after the port translation rule.

Note For applications that require application inspection for secondary channels (for example, FTP and VoIP), NAT automatically translates the secondary ports.

Following are some other uses of static NAT with port translation.

Static NAT with Identity Port Translation

You can simplify external access to internal resources. For example, if you have three separate servers that provide services on different ports (such as FTP, HTTP, and SMTP), you can give external users a single IP address to access those services. You can then configure static NAT with identity port translation to map the single external IP address to the correct IP addresses of the real servers based on the port they are trying to access. You do not need to change the port, because the servers are using the standard ones (21, 80, and 25 respectively).

Static NAT with Port Translation for Non-Standard Ports

You can also use static NAT with port translation to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect

to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell web users to connect to non-standard port 6785, and then undo translation to port 80.

Static Interface NAT with Port Translation

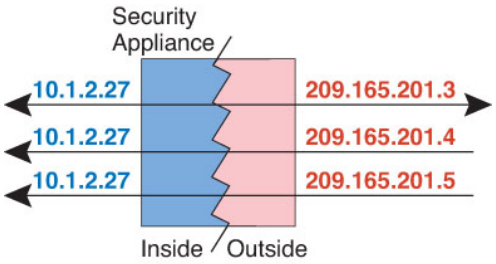
You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the device's outside interface to an inside host, then you can map the inside host IP address/port 23 to the outside interface address/port 23.

One-to-many static NAT

Typically, you configure static NAT with a one-to-one mapping. However, in some cases, you might want to configure a single real address to several mapped addresses (one-to-many). When you configure one-to-many static NAT, when the real host initiates traffic, it always uses the first mapped address. However, for traffic initiated to the host, you can initiate traffic to any of the mapped addresses, and they will be untranslated to the single real address.

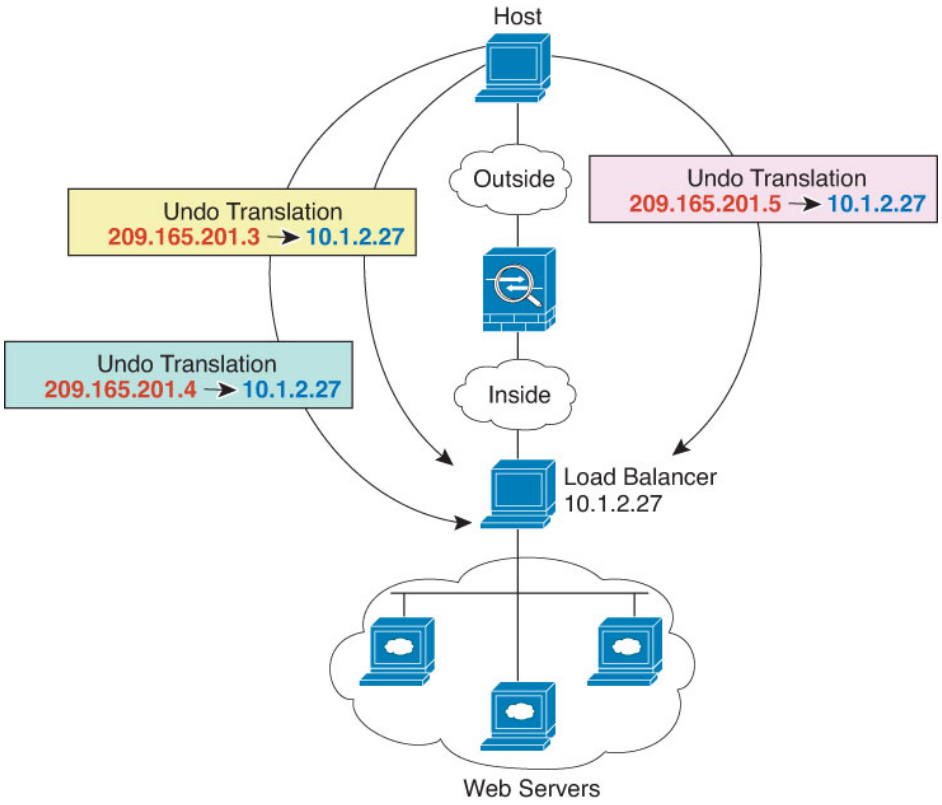
The following figure shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/first mapped IP is technically the only bidirectional translation.

Figure 353: One-to-Many Static NAT



For example, you have a load balancer at 10.1.2.27. Depending on the URL requested, it redirects traffic to the correct web server.

Figure 354: One-to-Many Static NAT Example



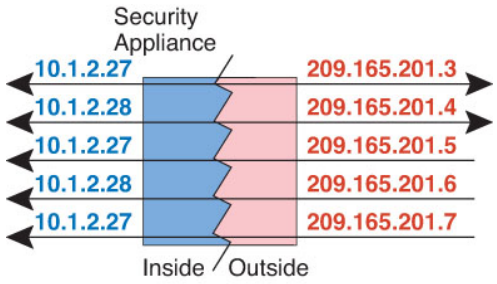
Other mapping scenarios (not recommended)

NAT has the flexibility to allow any kind of static mapping scenario: one-to-one, one-to-many, but also few-to-many, many-to-few, and many-to-one mappings. We recommend using only one-to-one or one-to-many mappings. These other mapping options might result in unintended consequences.

Functionally, few-to-many is the same as one-to-many; but because the configuration is more complicated and the actual mappings may not be obvious at a glance, we recommend creating a one-to-many configuration for each real address that requires it. For example, for a few-to-many scenario, the few real addresses are mapped to the many mapped addresses in order (A to 1, B to 2, C to 3). When all real addresses are mapped, the next mapped address is mapped to the first real address, and so on until all mapped addresses are mapped (A to 4, B to 5, C to 6). This results in multiple mapped addresses for each real address. Just like a one-to-many configuration, only the first mappings are bidirectional; subsequent mappings allow traffic to be initiated *to* the real host, but all traffic *from* the real host uses only the first mapped address for the source.

The following figure shows a typical few-to-many static NAT scenario.

Figure 355: Few-to-Many Static NAT

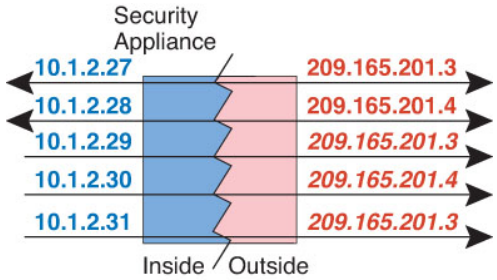


For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining higher real addresses can initiate traffic, but traffic cannot be initiated to them (returning traffic for a connection is directed to the correct real address because of the unique 5-tuple (source IP, destination IP, source port, destination port, protocol) for the connection).

Note Many-to-few or many-to-one NAT is not PAT. If two real hosts use the same source port number and go to the same outside server and the same TCP destination port, and both hosts are translated to the same IP address, then both connections will be reset because of an address conflict (the 5-tuple is not unique).

The following figure shows a typical many-to-few static NAT scenario.

Figure 356: Many-to-Few Static NAT



Instead of using a static rule this way, we suggest that you create a one-to-one rule for the traffic that needs bidirectional initiation, and then create a dynamic rule for the rest of your addresses.

Configure static auto NAT

Use static auto NAT rules to translate addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.

- **Translated Source**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
  - **Address**—Create a network object or group containing hosts, ranges, or subnets. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** drop-down button to create a new rule and select the **Auto NAT rule** option in the drop-down.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Static**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.

- (Optional.) **Original Port**, **Translated Port**—If you need to translate a TCP or UDP port, select the protocol in **Original Port**, and type the original and translated port numbers. For example, you can translate TCP/80 to 8080 if necessary.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 1677. This option is not available if you are doing port translation.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Net to Net Mapping**—For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure static manual NAT

Use static manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Static NAT translates addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.

- **Address**—Create a network object or group containing hosts, range, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule and select the **Manual NAT rule** option in the drop-down.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

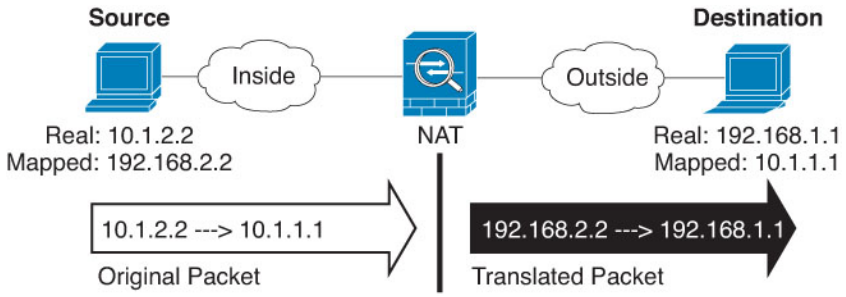
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

- **Original Source Port**, **Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port**, **Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 1677. This option is not available if you are doing port translation.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Net to Net Mapping**—For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Unidirectional**—Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Step 9 Click **Save** to add the rule.

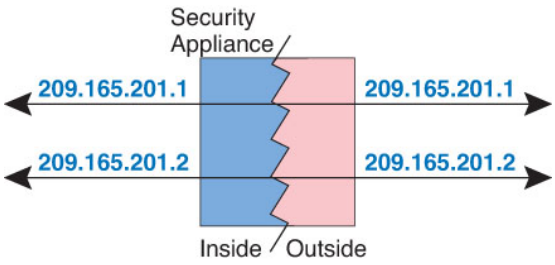
Step 10 Click **Save** on the NAT page to save your changes.

Identity NAT

You might have a NAT configuration in which you need to translate an IP address to itself. For example, if you create a broad rule that applies NAT to every network, but want to exclude one network from NAT, you can create a static NAT rule to translate that network to itself.

The following figure shows a typical identity NAT scenario.

Figure 357: Identity NAT



The following topics explain how to configure identity NAT.

Configure identity auto NAT

Use static identity auto NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—A network object or group with the exact same contents as the original source object. You can use the same object.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** drop-down button to create a new rule and select the **Auto NAT rule** option in the drop-down.
- Click **Edit** ( ) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Static**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Do not configure the **Original Port** and **Translated Port** options for identity NAT.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **IPv6**—Do not configure this option for identity NAT.
- **Net to Net Mapping**—Do not configure this option for identity NAT.

- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure identity manual NAT

Use static identity manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Use static identity NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—The same object or group as the original source. Optionally, you can select a different object that has the exact same contents.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports. You can use the same object for identity NAT.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule and select the **Manual NAT rule** option in the drop-down.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

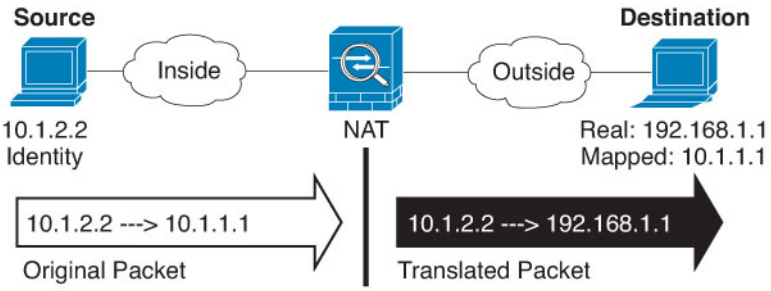
- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet where you perform identity NAT on the inside host but translate the outside host.



- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Interface Object** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—The same object or group as the original source. Optionally, you can select a different object that has the exact same contents.

- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

- **Original Source Port**, **Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port**, **Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.
- **Unidirectional**—Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

NAT rule properties for Firewall Threat Defense

Use network address translation (NAT) rules to translate IP addresses to other IP addresses. You would typically use NAT rules to convert private addresses to publicly routable addresses. The translation can be from one address to another, or you can use port address translation (PAT) to translate many addresses to one or a few addresses, using port numbers to distinguish among the source addresses.

NAT rules include the following basic properties. The properties are the same for auto NAT and manual NAT rules except where indicated.

NAT Type

Whether you want to configure a **Manual NAT Rule** or an **Auto NAT Rule**. Auto NAT translates the source address only, and you cannot make different translations based on the destination address. Because auto NAT is more simple to configure, use it unless you need the added features of manual NAT. For more information on the differences, see Auto NAT and Manual NAT, on page 1571.

Type

Whether the translation rule is **Dynamic** or **Static**. Dynamic translation automatically chooses the mapped address from a pool of addresses, or an address/port combination when implementing PAT. Use static translation if you want to precisely define the mapped address/port.

Enable (Manual NAT only.)

Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page. You cannot disable auto NAT rules.

Insert (Manual NAT only.)

Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Description (Optional. Manual NAT only.)

A description of the purpose of the rule.

The following topics describe the tabs for the NAT rules properties.

Interface objects NAT properties

Interface objects (security zones or interface groups) define the interfaces to which a NAT rule applies. In routed mode, you can use the default "any" for both source and destination to apply to all interfaces of all assigned devices. However, you typically want to select specific source and destination interfaces.

Notes

- The concept of “any” interface does not apply to bridge group member interfaces. When you specify “any” interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. You cannot configure NAT for the bridge virtual interface (BVI) itself, you can configure NAT for member interfaces only.

  If you select interface objects, a NAT rule will be configured on an assigned device only if the device has interfaces included in all selected objects. For example, if you select both source and destination security zones, both zones must contain one or more interface for a given device.

- If more than one interface in an interface object exists on a given device, identical rules are created for each interface. This can become an issue for static NAT rules that include destination translation. Because NAT rules are applied based on first hit rule, only the rule created for the first interface configured for the object matches traffic. When configuring static NAT with destination translation, use interface objects that include at most one interface per device assigned to the NAT policy to ensure you are getting the desired results.

Source Interface Objects, Destination Interface Objects

(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped

interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Translation properties for auto NAT

Use the **Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to auto NAT only.

Original Source (Always required.)

The network object that contains the addresses you are translating. This must be a network object (not a group), and it can be a host, range, or subnet.

You cannot create auto NAT rules for the system-defined any-ipv4 or any-ipv6 objects.

Translated Source (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.
- **Dynamic PAT**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Do not configure a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Do not configure a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Select the PAT pool object on **PAT Pool**.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. The object or group can contain hosts, ranges, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on the **Advanced** tab. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Port, Translated Port (Static NAT only.)

If you need to translate a TCP or UDP port, select the protocol in **Original Port**, and type the original and translated port numbers. For example, you can translate TCP/80 to 8080 if necessary. Do not configure these options for identity NAT.

Translation properties for manual NAT

Use the **Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to manual NAT only. All are optional except as indicated.

Original Source (Always required.)

The network object or group that contains the addresses you are translating. This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can specify **Any** in the rule.

Translated Source (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback. If the object contains one host address only, it is used for PAT.
- **Dynamic PAT**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Do not configure a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Do not configure a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Select the PAT pool object on **PAT Pool**.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. The object or group can contain hosts, ranges, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on the **Advanced** tab. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Destination

The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Translated Destination

The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

You can use a network object that specifies a fully-qualified domain name as the translated destination; for more information, see FQDN destination guidelines, on page 1580.

Original Source Port, Translated Source Port, Original Destination Port, Translated Destination Port

The port objects that define the source and destination services for the original and translated packets. You can translate the ports, or select the same object to make the rule sensitive to the service without translating the ports. Keep the following rules in mind when configuring services:

- (Dynamic NAT or PAT.) You cannot do translation on the **Original Source Port** and **Translated Source Port**. You can do translation on the destination port only.
- NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

PAT pool NAT properties

When you configure dynamic NAT, you can define a pool of addresses to use for port address translation using the properties on the **PAT Pool** tab.

Enable PAT Pool

Select this option to configure a pool of addresses for PAT.

PAT

The addresses to use for the PAT pool, one of the following:

- **Address**—The object that defines the PAT pool addresses, either a network object that includes a range, or a network object group that contains hosts, ranges, or both. You cannot include subnets. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- **Destination Interface IP**—Indicates that you want to use the destination interface as the PAT address. For this option, you must select a specific **Destination Interface Object**; you cannot use **Any** as the destination interface. This is another way to implement interface PAT.

Round Robin

To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one

address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.

Extended PAT Table

To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.

Flat Port Range; Include Reserved Ports

To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option. For the Firewall Threat Defense devices running version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.

Block Allocation

To enable port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation is compatible with round robin, but you cannot use it with the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Advanced NAT properties

When you configure NAT, you can configure properties that provide specialized services in the **Advanced** options. All of these properties are optional: configure them only if you need the service.

Translate DNS replies that match this rule

Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see DNS query and response rewriting using NAT, on page 1677. This option is not available if you are doing port translation in a static NAT rule.

Fallthrough to Interface PAT (Destination Interface) (Dynamic NAT only.)

Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option. You cannot select this option if you already configured interface PAT as the translated address. You also cannot select the option if you configure a PAT pool.

IPv6

Whether to use the IPv6 address of the destination interface for interface PAT.

Net to Net Mapping (Static NAT only.)

For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.

Do not proxy ARP on Destination Interface (Static NAT only.)

Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Perform Route Lookup for Destination Interface (Static Identity NAT only. Routed mode only.)

If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Unidirectional (Manual NAT only, static NAT only.)

Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Translating IPv6 networks

In cases where you need to pass traffic between IPv6-only and IPv4-only networks, you need to use NAT to convert between the address types. Even with two IPv6 networks, you might want to hide internal addresses from the outside network.

You can use the following translation types with IPv6 networks:

- NAT64, NAT46—Translates IPv6 packets into IPv4 and vice versa. You need to define two policies, one for the IPv6 to IPv4 translation, and one for the IPv4 to IPv6 translation. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.



Note NAT46 supports static mappings only.

- NAT66—Translates IPv6 packets to a different IPv6 address. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.

Note NAT64 and NAT 46 are possible on standard routed interfaces only. NAT66 is possible on both routed and bridge group member interfaces.

NAT64/46: translating IPv6 addresses to IPv4

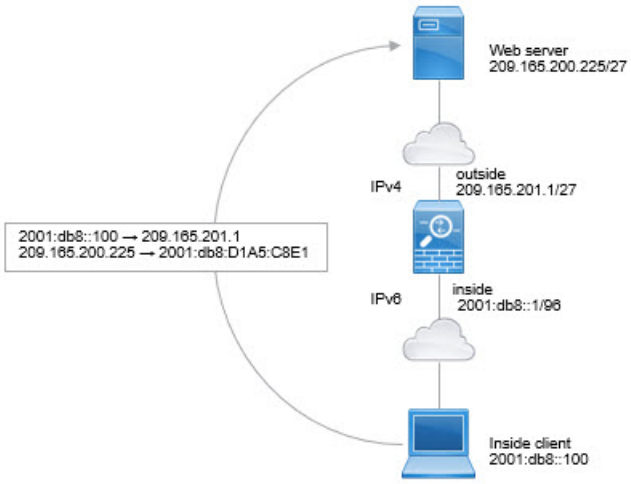
When traffic goes from an IPv6 network to an IPv4-only network, you need to convert the IPv6 address to IPv4, and return traffic from IPv4 to IPv6. You need to define two address pools, an IPv4 address pool to bind IPv6 addresses in the IPv4 network, and an IPv6 address pool to bind IPv4 addresses in the IPv6 network.

- The IPv4 address pool for the NAT64 rule is normally small and typically might not have enough addresses to map one-to-one with the IPv6 client addresses. Dynamic PAT might more easily meet the possible large number of IPv6 client addresses compared to dynamic or static NAT.
- The IPv6 address pool for the NAT46 rule can be equal to or larger than the number of IPv4 addresses to be mapped. This allows each IPv4 address to be mapped to a different IPv6 address. NAT46 supports static mappings only, so you cannot use dynamic PAT.

You need to define two policies, one for the source IPv6 network, and one for the destination IPv4 network. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.

NAT64/46 example: inside IPv6 network with outside IPv4 internet

Following is a straight-forward example where you have an inside IPv6-only network, and you want to convert to IPv4 for traffic sent to the Internet. This example assumes you do not need DNS translation, so you can perform both the NAT64 and NAT46 translations in a single manual NAT rule.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network.

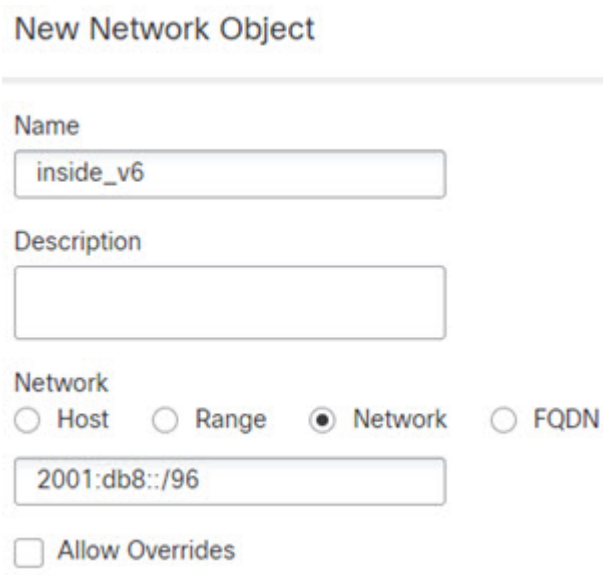
Procedure

Step 1 Create the network object that defines the inside IPv6 network.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8::/96.



c) Click **Save**.

Step 2 Create the manual NAT rule to translate the IPv6 network to IPv4 and back again.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

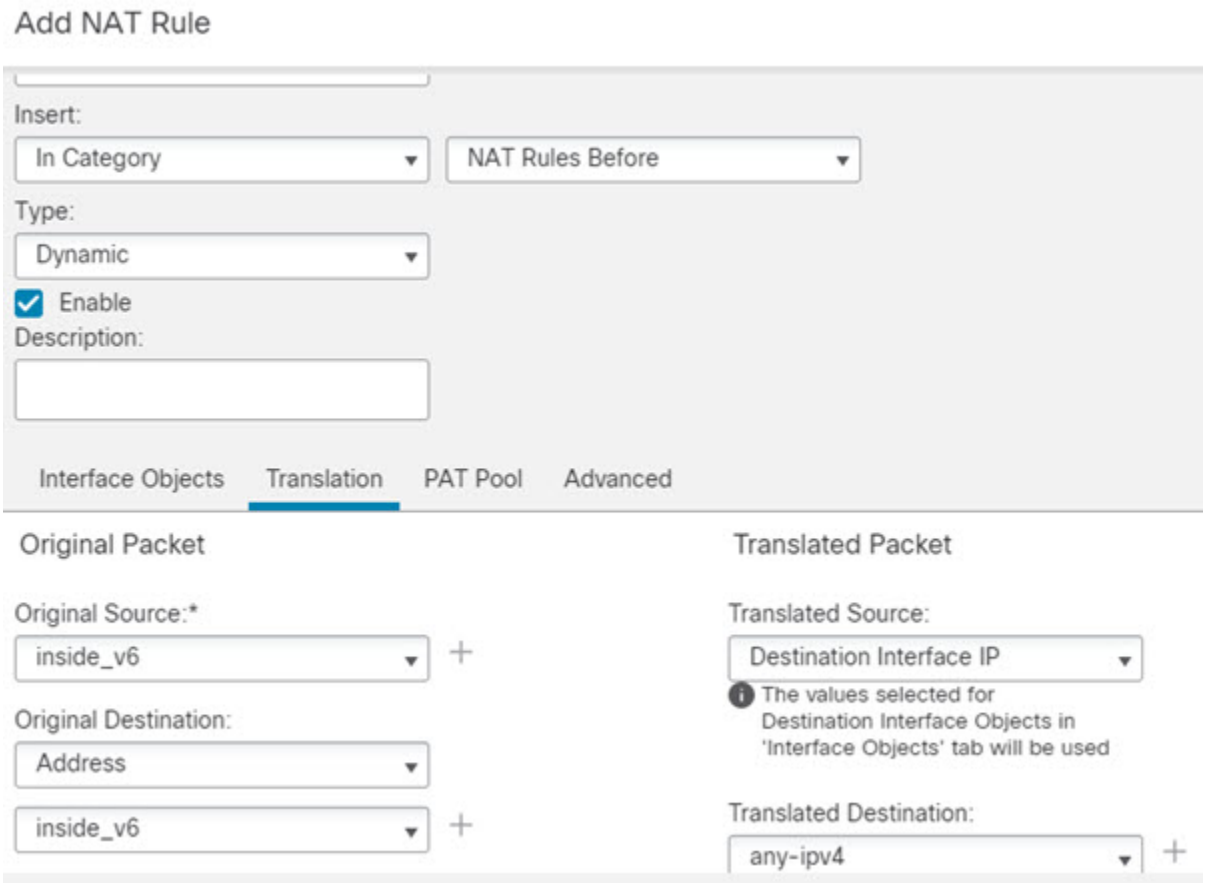
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** = **Destination Interface IP**.
- **Original Destination** = inside_v6 network object.
- **Translated Destination** = any-ipv4 network object.



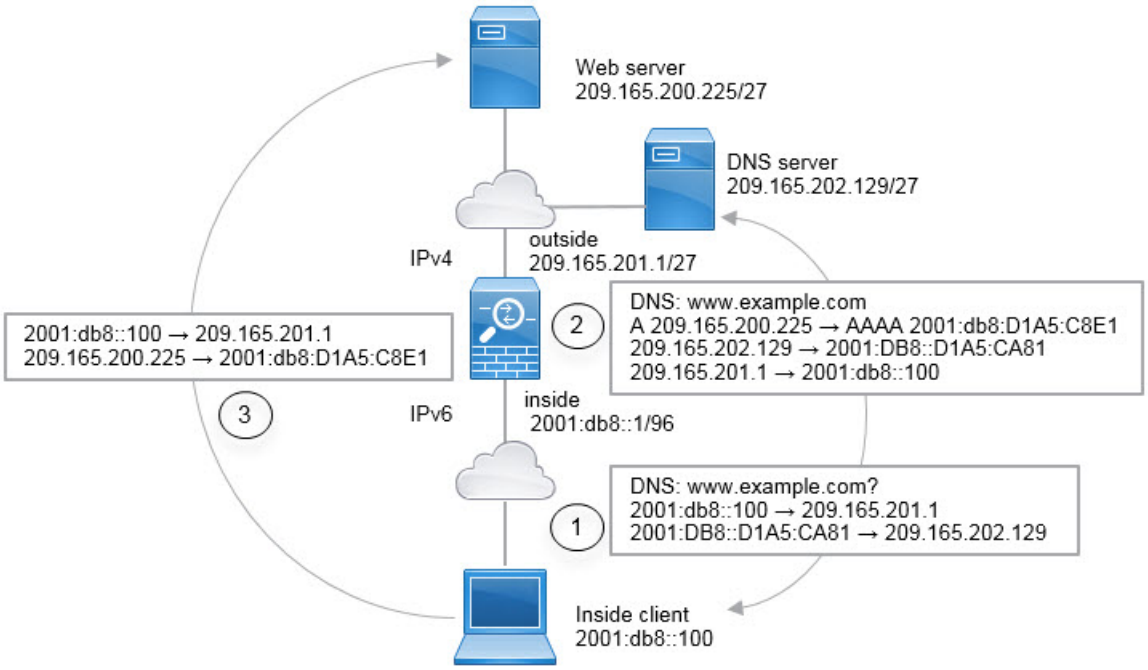
f) Click **OK**.

With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface. Conversely, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method.

g) Click **Save** on the NAT rules page.

NAT64/46 example: inside IPv6 network with outside IPv4 internet and DNS translation

Following is a typical example where you have an inside IPv6-only network, but there are some IPv4-only services on the outside Internet that internal users need.



In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network. You enable DNS rewrite on the NAT46 rule, so that replies from the external DNS server can be converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

Following is a typical sequence for a web request where a client at 2001:DB8::100 on the internal IPv6 network tries to open www.example.com.

1. The client's computer sends a DNS request to the DNS server at 2001:DB8::D1A5:CA81. The NAT rules make the following translations to the source and destination in the DNS request:
   - 2001:DB8::100 to a unique port on 209.165.201.1 (The NAT64 interface PAT rule.)
   - 2001:DB8::D1A5:CA81 to 209.165.202.129 (The NAT46 rule. D1A5:CA81 is the IPv6 equivalent of 209.165.202.129.)
2. The DNS server responds with an A record indicating that www.example.com is at 209.165.200.225. The NAT46 rule, with DNS rewrite enabled, converts the A record to the IPv6-equivalent AAAA record, and translates 209.165.200.225 to 2001:db8:D1A5:C8E1in the AAAA record. In addition, the source and destination addresses in the DNS response are untranslated:
   - 209.165.202.129 to 2001:DB8::D1A5:CA81
   - 209.165.201.1 to 2001:db8::100
3. The IPv6 client now has the IP address of the web server, and makes an HTTP request to www.example.com at 2001:db8:D1A5:C8E1. (D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225.) The source and destination of the HTTP request are translated:
   - 2001:DB8::100 to a unique port on 209.156.101.54 (The NAT64 interface PAT rule.)
   - 2001:db8:D1A5:C8E1 to 209.165.200.225 (The NAT46 rule.)

The following procedure explains how to configure this example.

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

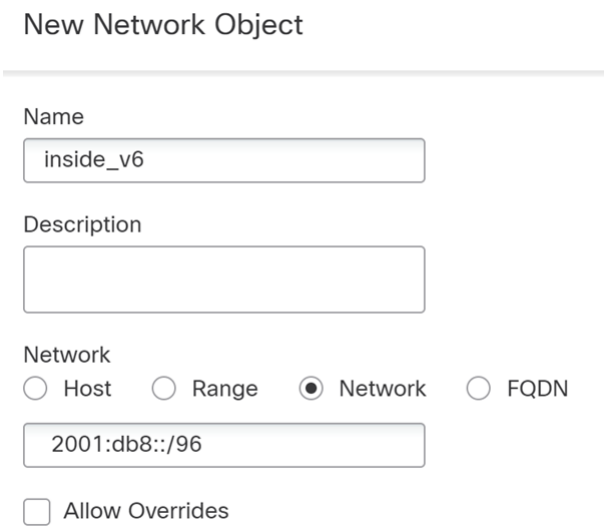
Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv4 networks.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the inside IPv6 network.

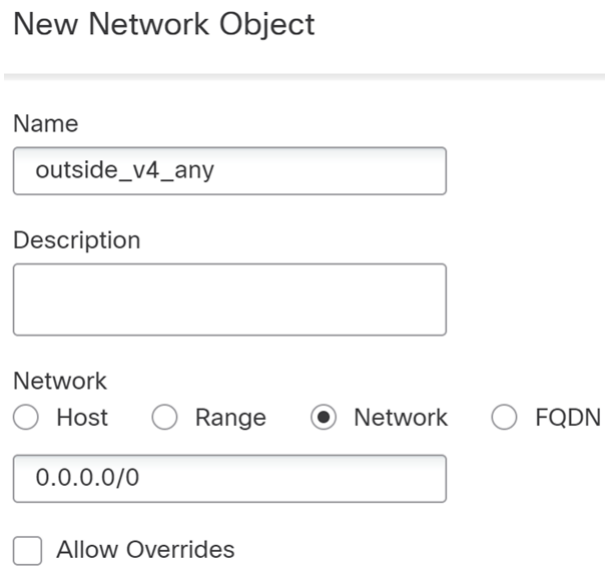
Name the network object (for example, inside_v6) and enter the network address, 2001:db8::/96.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the outside IPv4 network.

Name the network object (for example, outside_v4_any) and enter the network address 0.0.0.0/0.



e) Click **Save**.

Step 2 Configure the NAT64 dynamic PAT rule for the inside IPv6 network.

Step 3 Configure the static NAT46 rule for the outside IPv4 network.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

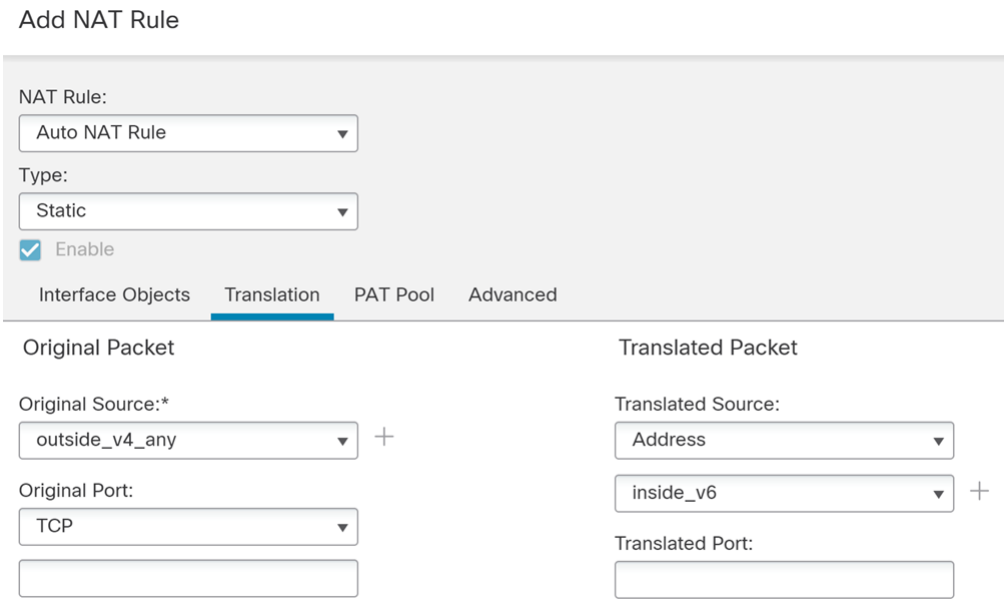
c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = outside_v4_any network object.
- **Translated Source** > **Address** = inside_v6 network object.

e) On **Advanced**, select **Translate DNS replies that match this rule**.



f) Click **OK**.

With this rule, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method. In addition, DNS responses are converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

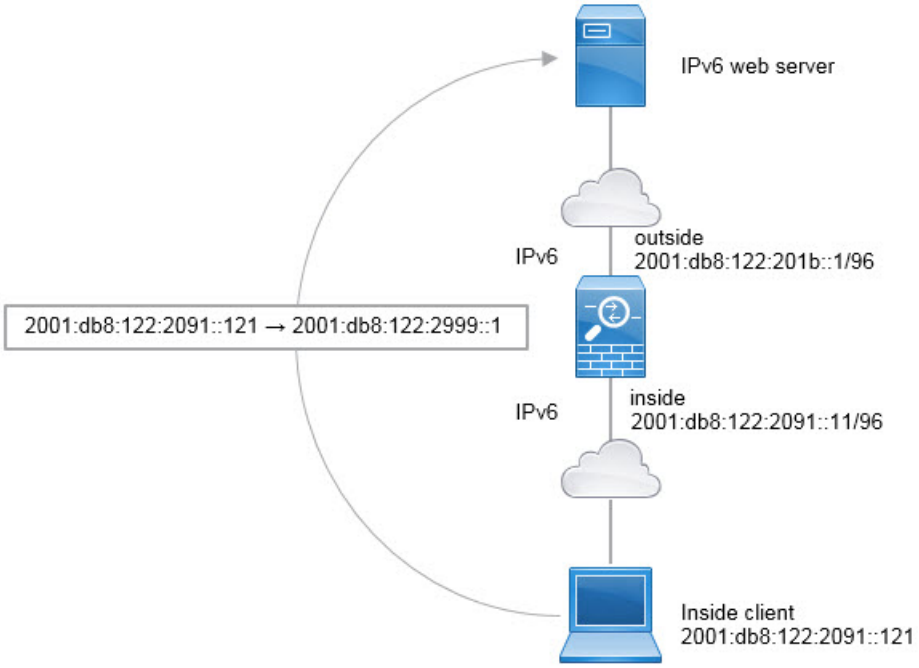
NAT66: translating IPv6 addresses to different IPv6 addresses

When going from an IPv6 network to another IPv6 network, you can translate the addresses to different IPv6 addresses on the outside network. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.

Because you are not translating between different address types, you need a single rule for NAT66 translations. You can easily model these rules using auto NAT. However, if you do not want to allow returning traffic, you can make the static NAT rule unidirectional using manual NAT only.

NAT66 example, static translation between networks

You can configure a static translation between IPv6 address pools using auto NAT. The following example explains how to convert inside addresses on the 2001:db8:122:2091::/96 network to outside addresses on the 2001:db8:122:2999::/96 network.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

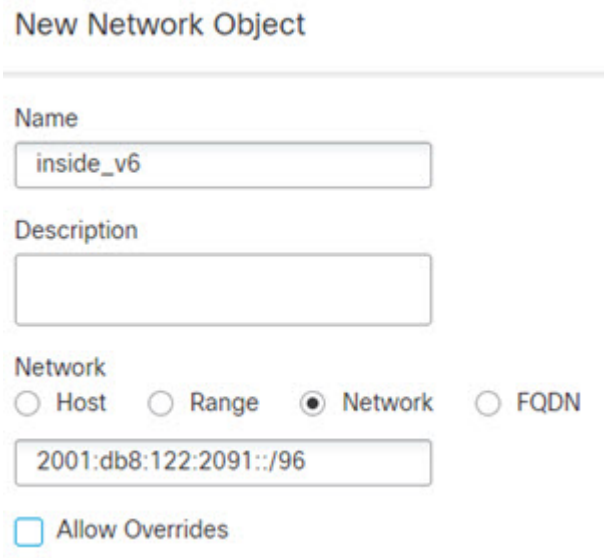
Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv6 NAT networks.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the inside IPv6 network.

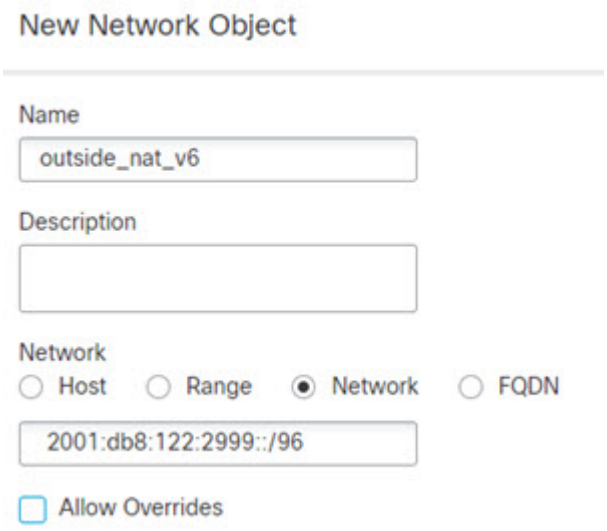
Name the network object (for example, inside_v6) and enter the network address, 2001:db8:122:2091::/96.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the outside IPv6 NAT network.

Name the network object (for example, outside_nat_v6) and enter the network address 2001:db8:122:2999::/96.



e) Click **Save**.

Step 2 Configure the static NAT rule for the inside IPv6 network.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

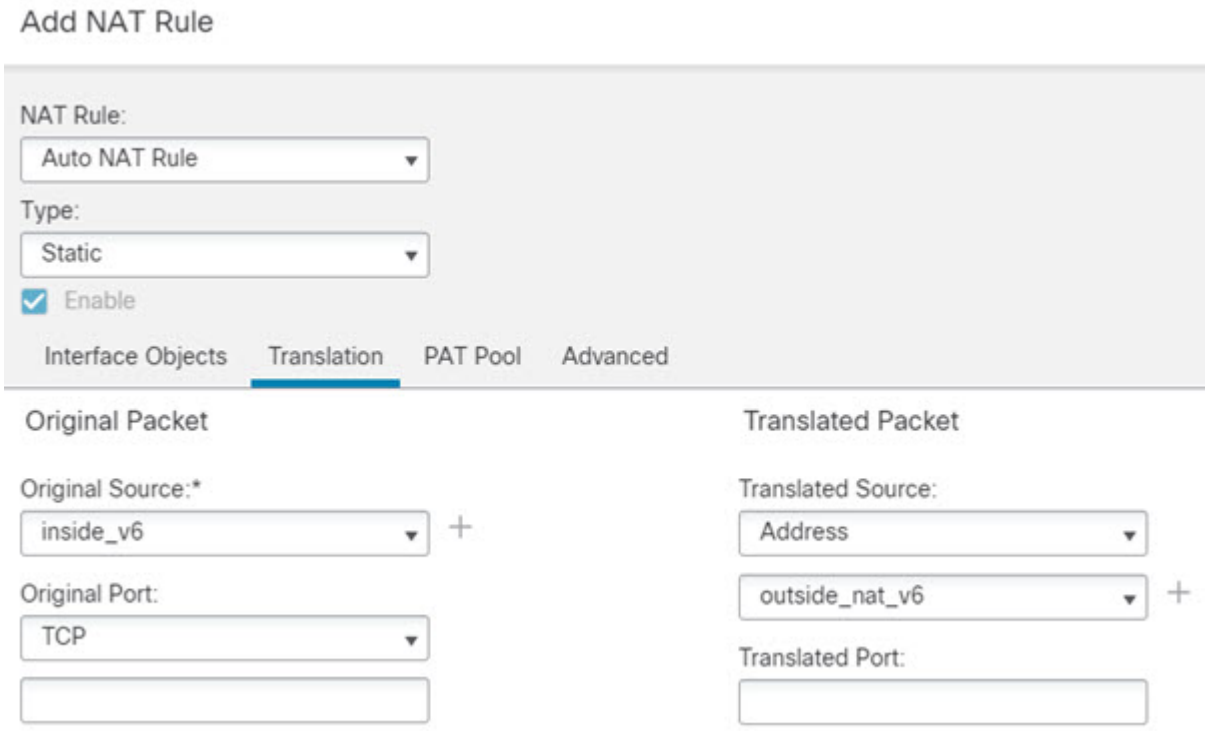
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** > **Address** = outside_nat_v6 network object.



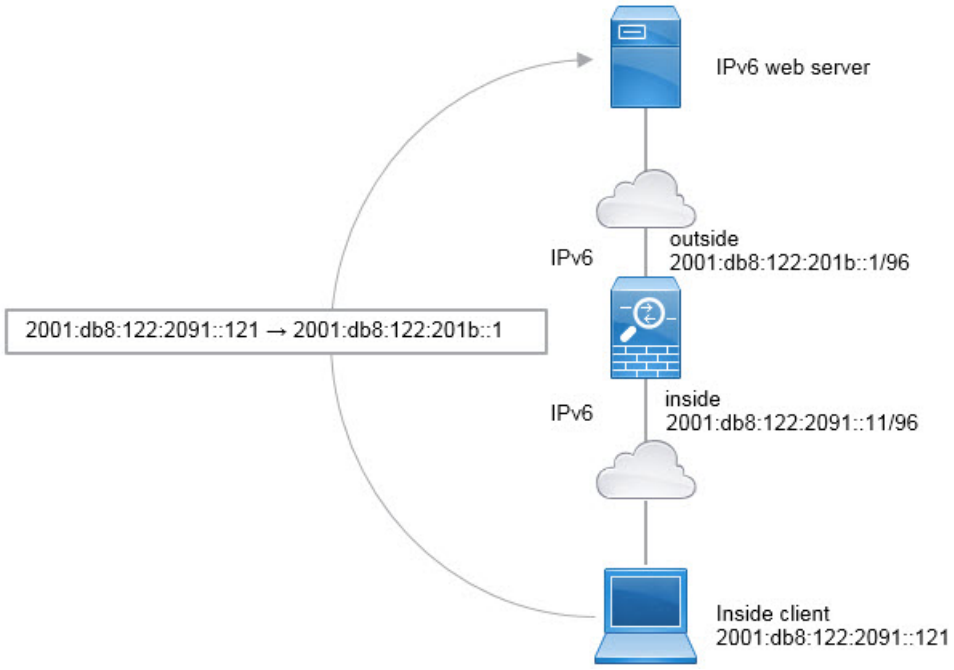
f) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a static NAT66 translation to an address on the 2001:db8:122:2999::/96 network.

NAT66 example, simple IPv6 interface PAT

A simple approach for implementing NAT66 is to dynamically assign internal addresses to different ports on the outside interface IPv6 address.

When you configure an interface PAT rule for NAT66, all the global addresses that are configured on that interface are used for PAT mapping. Link-local or site-local addresses for the interface are not used for PAT.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

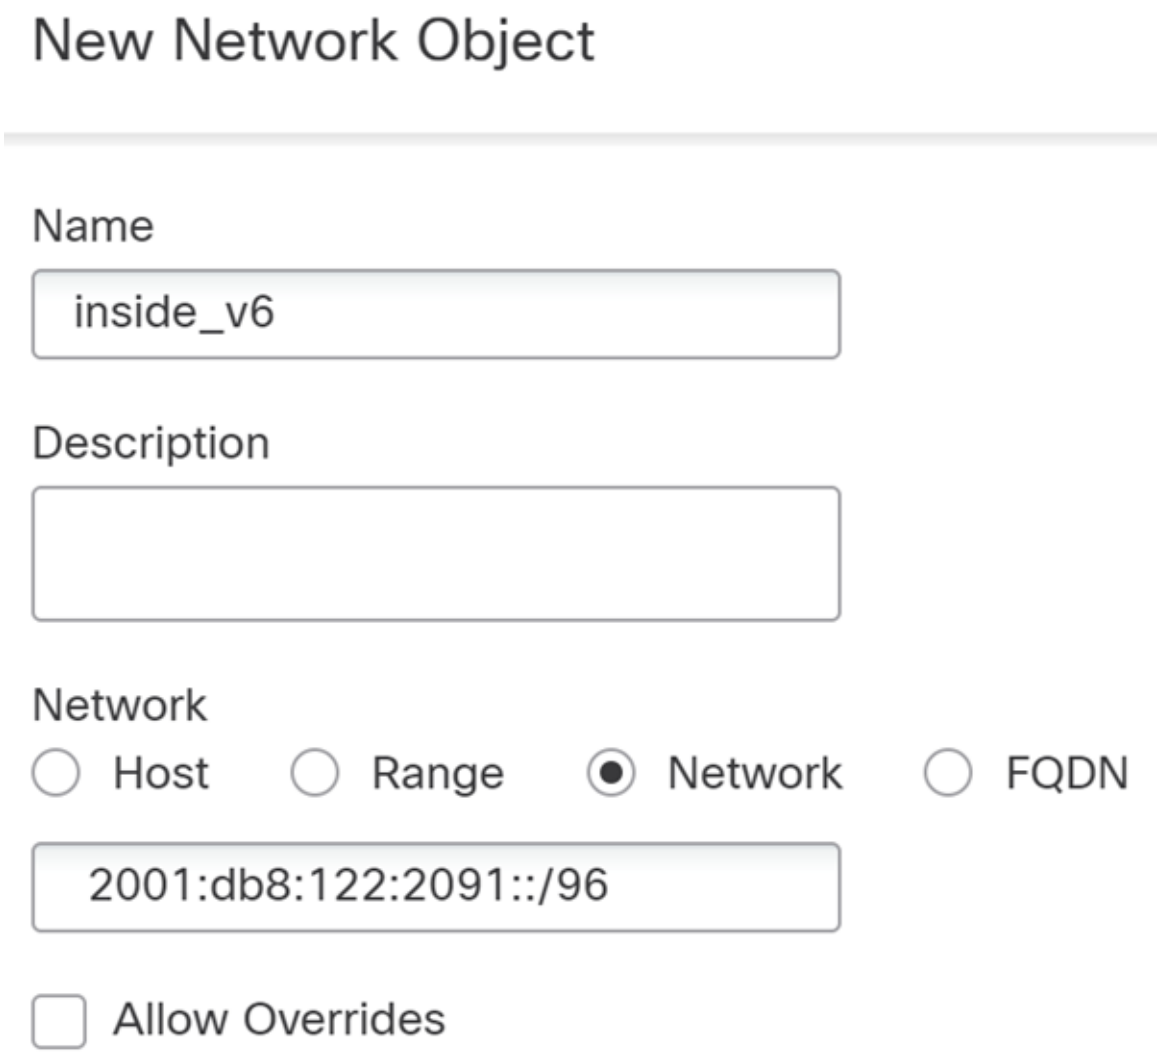
Procedure

Step 1 Create the network object that defines the inside IPv6 network.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8:122:2091::/96.



c) Click **Save**.

Step 2 Configure the dynamic PAT rule for the inside IPv6 network.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

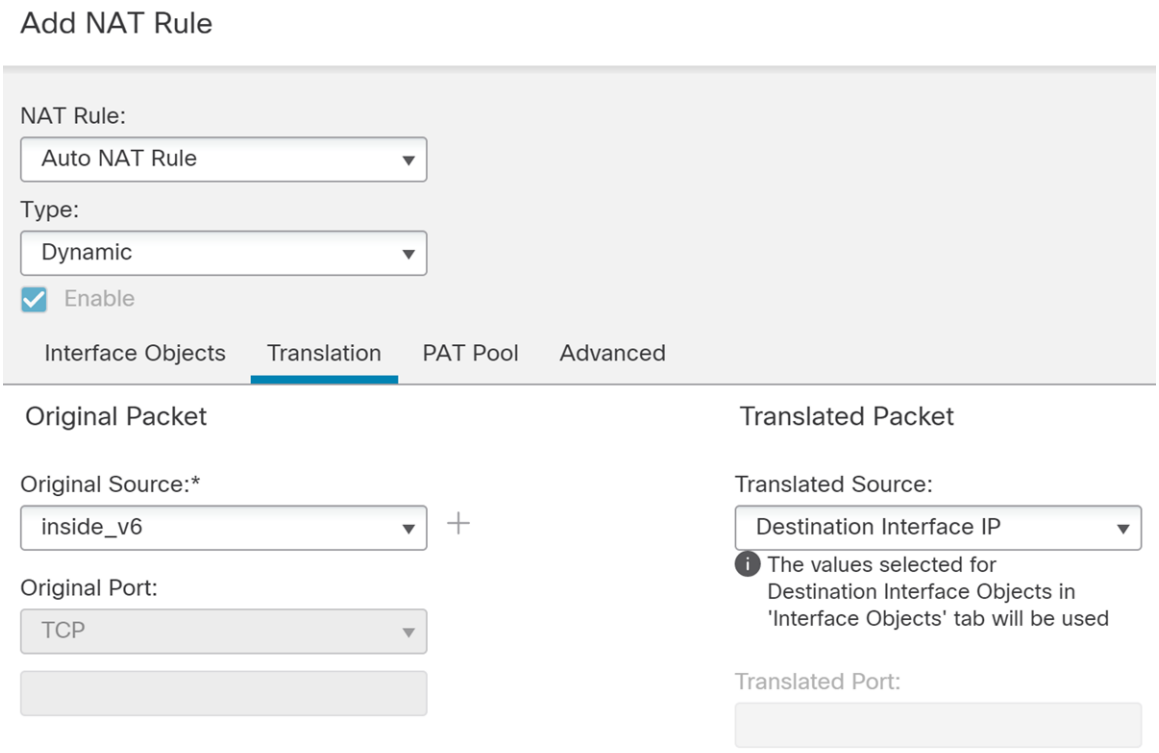
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** = **Destination Interface IP**.

f) On **Advanced**, select **IPv6**, which indicates that the IPv6 address of the destination interface should be used.



g) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a NAT66 PAT translation to one of the IPv6 global addresses configured for the outside interface.

Monitoring NAT

To monitor and troubleshoot NAT connections, log into the device CLI and use the following commands.

- **show nat** displays the NAT rules and per-rule hit counts. There are additional keywords to show other aspects of NAT.
- **show xlate** displays the actual NAT translations that are currently active.
- **clear xlate** lets you remove an active NAT translation. You might need to remove active translations if you alter NAT rules, because existing connections continue to use the old translation slot until the

connection ends. Clearing a translation allows the system to build a new translation for a client on the client's next connection attempt based on your new rules.

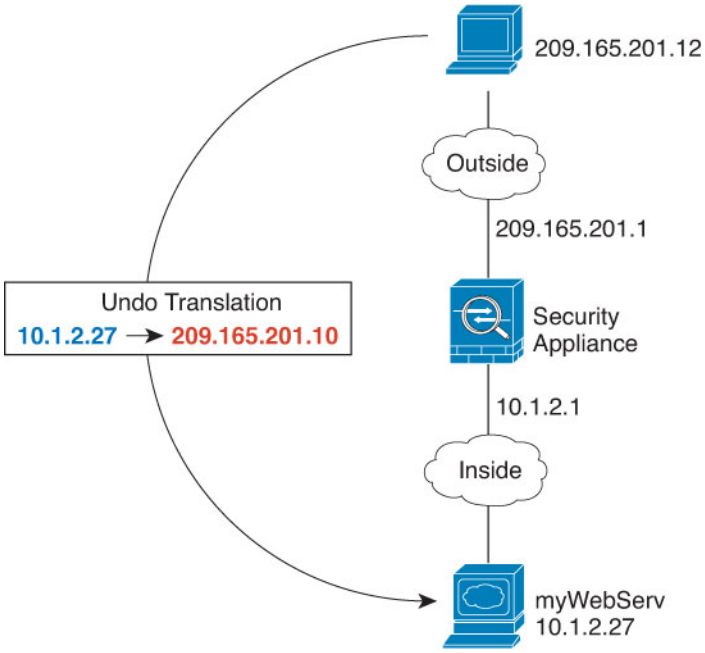
Examples for NAT

The following topics provide examples for configuring NAT.

Providing access to an inside web server (static auto NAT)

This example performs static NAT for an inside web server. The real address is on a private network, so a public address is required. Static NAT is necessary so hosts can initiate traffic to the web server at a fixed address.

Figure 358: Static NAT for an inside web server



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

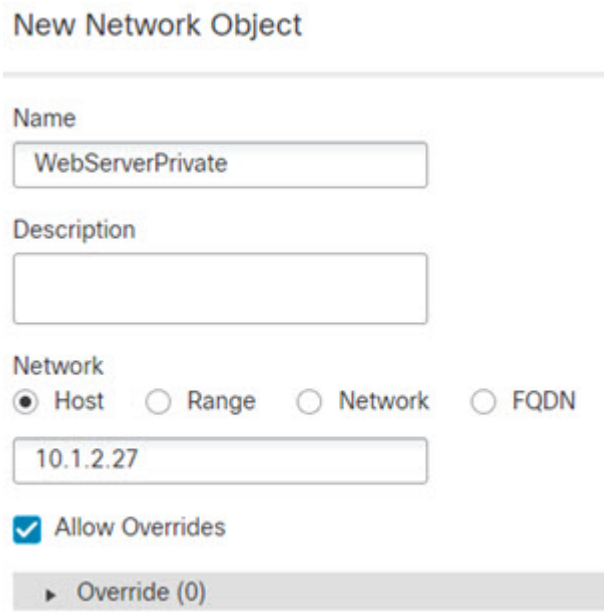
Procedure

Step 1 Create the network objects that define the server's private and public host addresses.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the web server's private address.

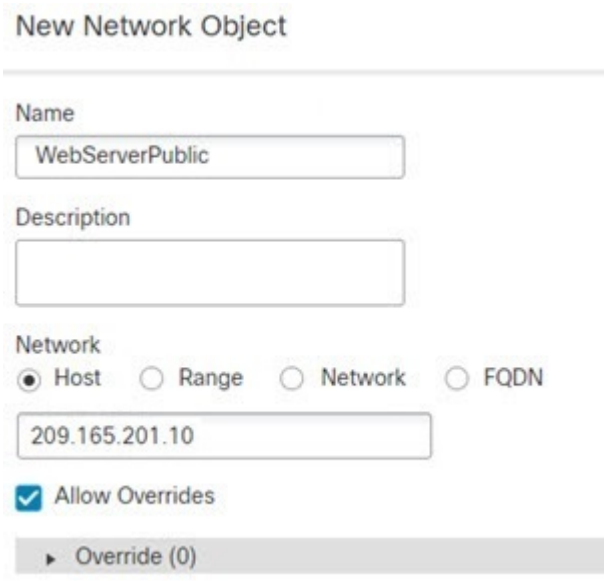
Name the network object (for example, WebServerPrivate) and enter the real host IP address, 10.1.2.27.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the public address.

Name the network object (for example, WebServerPublic) and enter the host address 209.165.201.10.



e) Click **Save**.

Step 2 Configure static NAT for the object.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

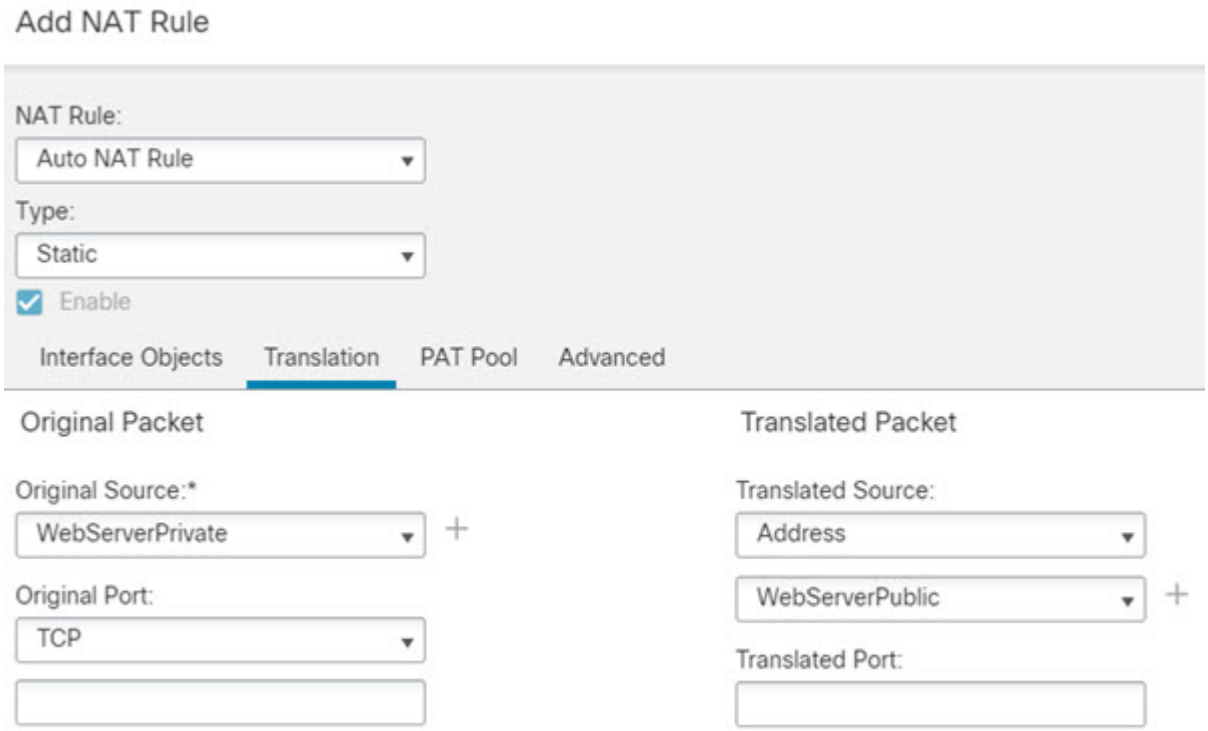
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = WebServerPrivate network object.
- **Translated Source** > **Address** = WebServerPublic network object.



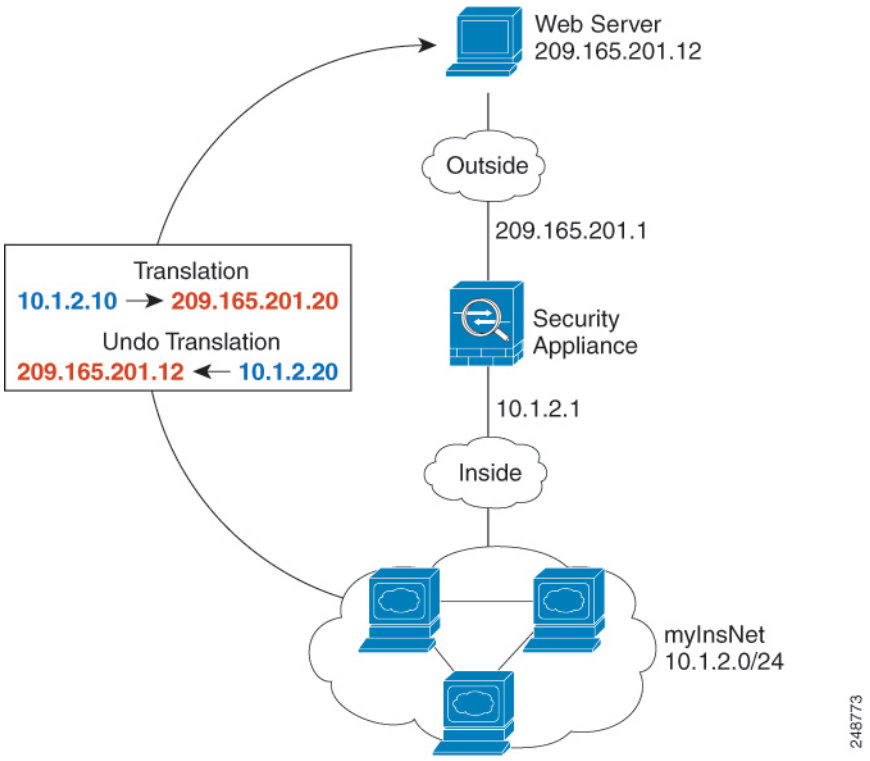
f) Click **Save**.

Step 3 Click **Save** on the NAT rule page.

Dynamic auto NAT for inside hosts and static NAT for an outside web server

The following example configures dynamic NAT for inside users on a private network when they access the outside. Also, when inside users connect to an outside web server, that web server address is translated to an address that appears to be on the inside network.

Figure 359: Dynamic NAT for Inside, Static NAT for Outside Web Server



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select and click **Add Network** > **Add ObjectObjects** > **Object Management** > **Interface**.

Procedure

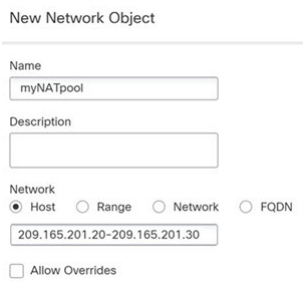
Step 1 Create a network object for the dynamic NAT pool to which you want to translate the inside addresses.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the dynamic NAT pool.

Name the network object (for example, myNATpool) and enter the network range 209.165.201.20-209.165.201.30.

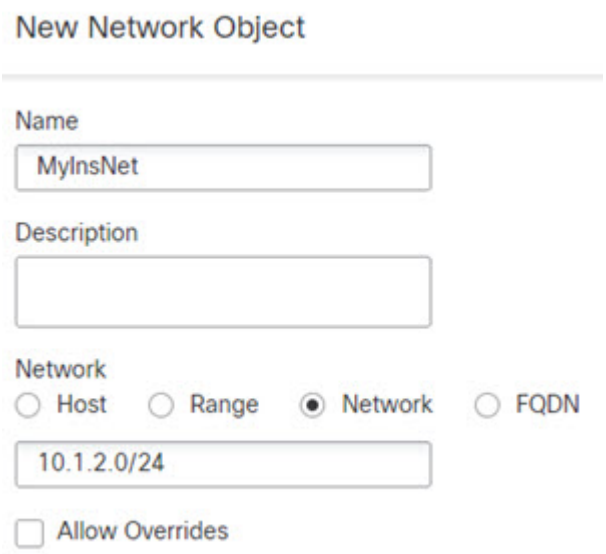


d) Click **Save**.

Step 2 Create a network object for the inside network.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, MyInsNet) and enter the network address 10.1.2.0/24.

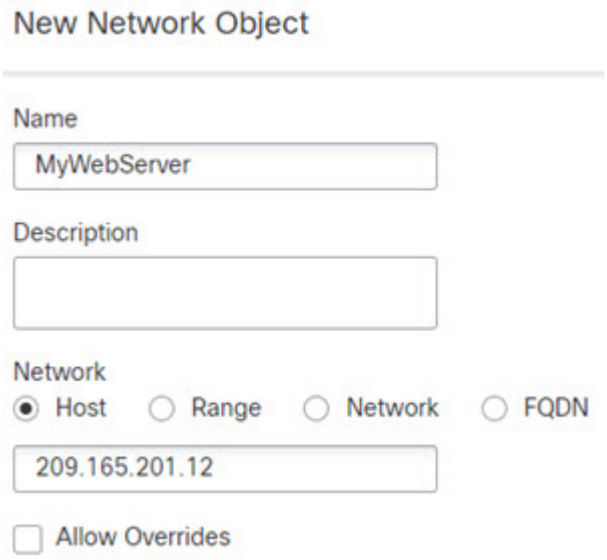


c) Click **Save**.

Step 3 Create a network object for the outside web server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, MyWebServer) and enter the host address 209.165.201.12.

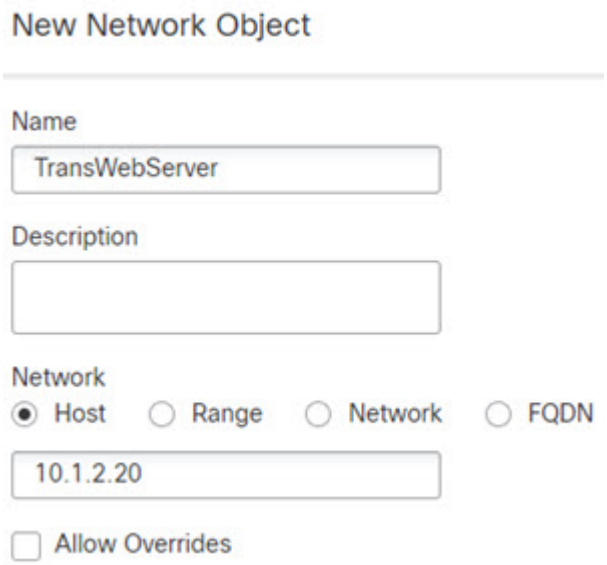


c) Click **Save**.

Step 4 Create a network object for the translated web server address.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, TransWebServer) and enter the host address 10.1.2.20.



c) Click **Save**.

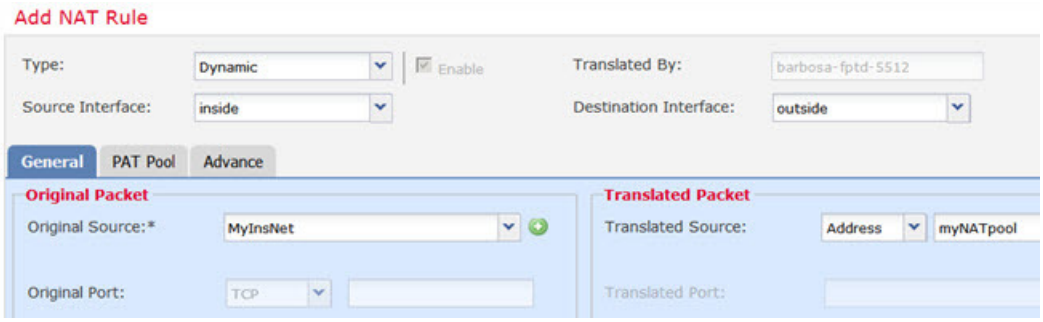
Step 5 Configure dynamic NAT for the inside network using the dynamic NAT pool object.

a) Select **Devices** > **Device Management** and edit the Threat Defense device on which you are defining the rule. Select **NAT**.

b) Click **Add NAT Rule** > **Add Auto NAT**.

c) Configure the following properties:

- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Source** = myInsNet network object.
- **Translated Source** > **Address**= myNATpool network group.



d) Click **Save**.

Step 6 Configure dynamic NAT for the inside network using the dynamic NAT pool object.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:

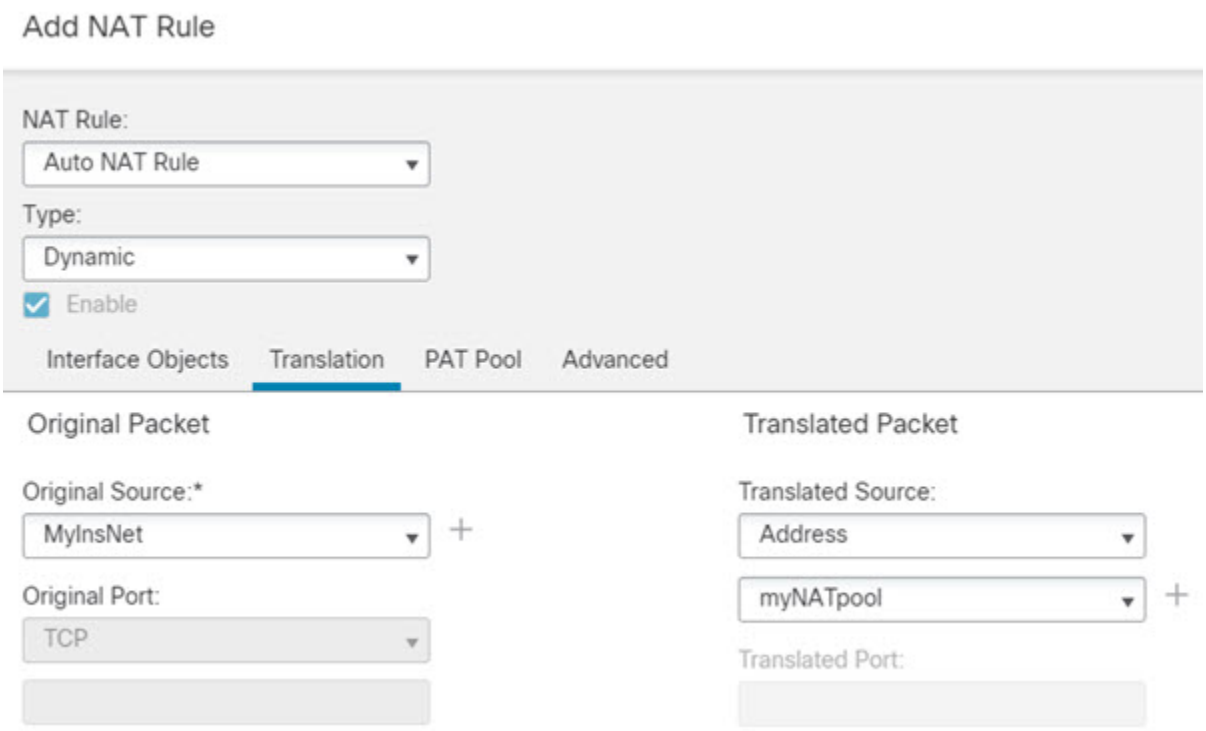
- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = myInsNet network object.
- **Translated Source** > **Address**= myNATpool network group.



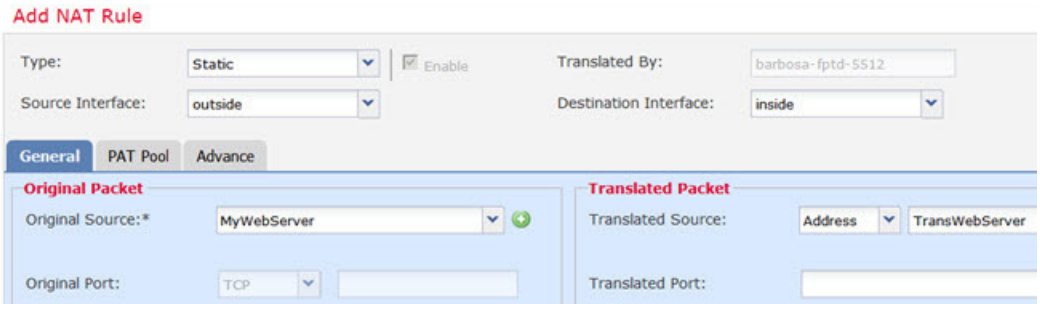
f) Click **Save**.

Step 7 Configure static NAT for the web server.

a) Click **Add NAT Rule** > **Add Auto NAT**.

b) Configure the following properties:

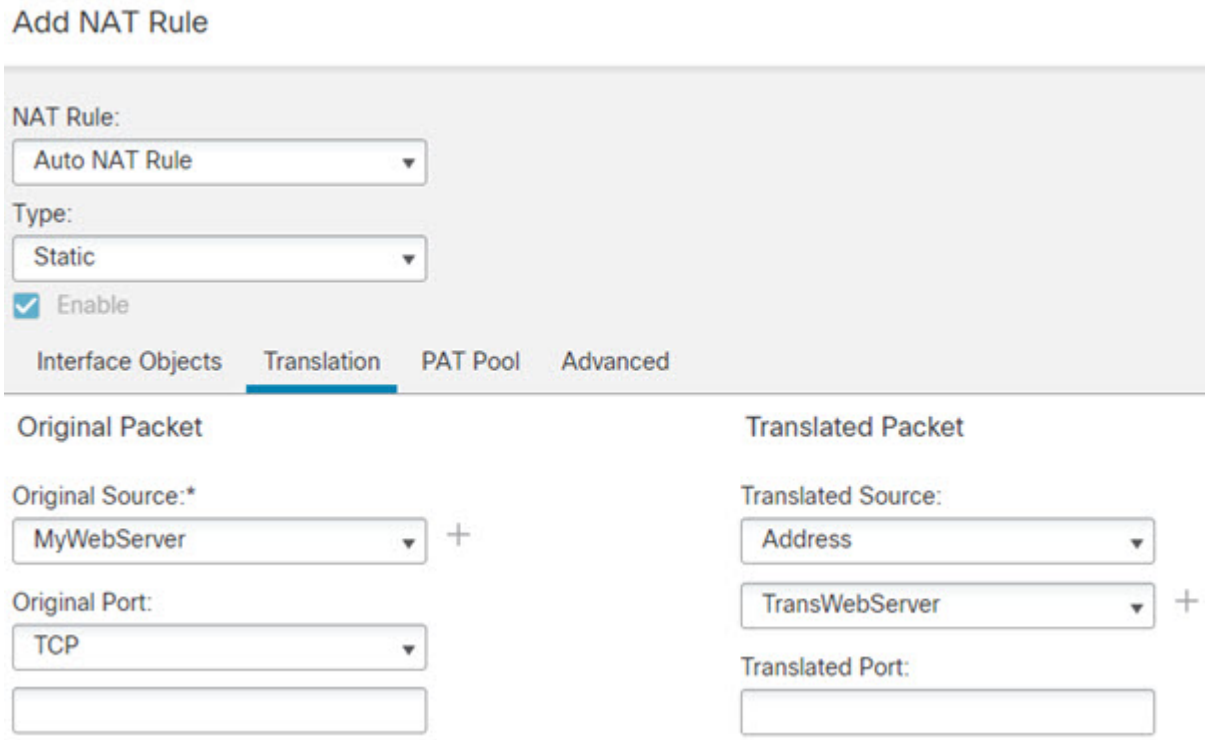
- **Type** = Static.
- **Source Interface** = outside.
- **Destination Interface** = inside.
- **Original Source** = myWebServer network object.
- **Translated Source** > **Address**= TransWebServer network object.



c) Click **Save**.

Step 8 Configure static NAT for the web server.

a) Click **Add Rule**.
b) Configure the following properties:
   - **NAT Rule** = Auto NAT Rule.
   - **Type** = Static.
c) On **Interface Objects**, configure the following:
   - **Source Interface Objects** = outside.
   - **Destination Interface Objects** = inside.
d) On **Translation**, configure the following:
   - **Original Source** = myWebServer network object.
   - **Translated Source** > **Address**= TransWebServer network object.



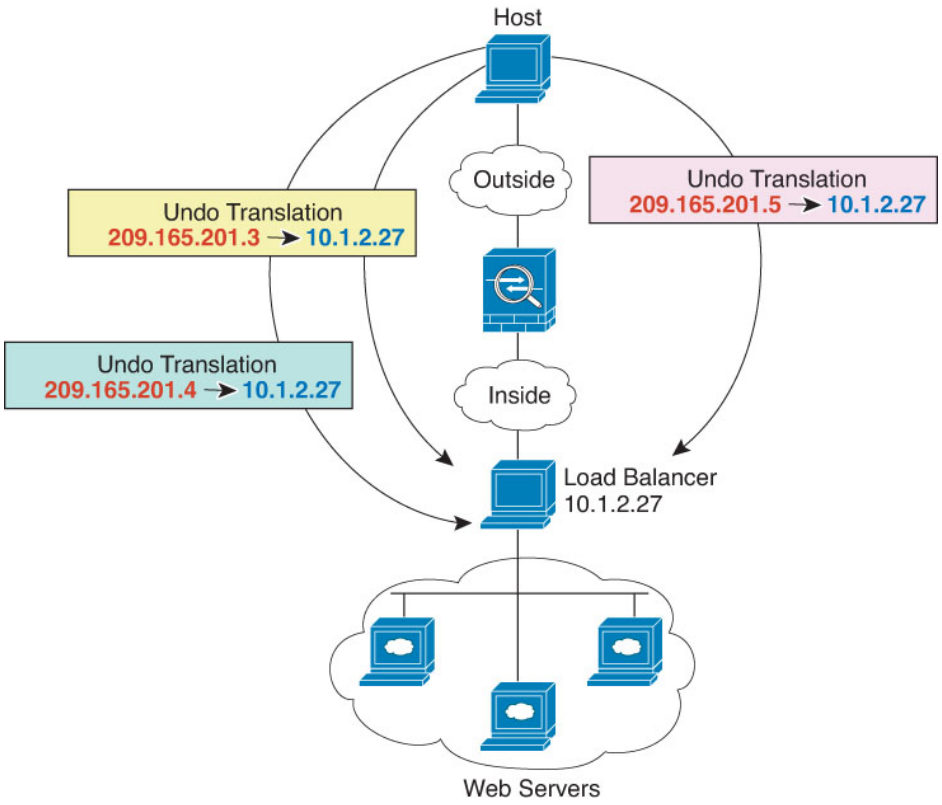
e) Click **Save**.

Step 9 Click **Save** on the NAT rule page.

Configure inside load balancer with multiple mapped addresses

The following example shows an inside load balancer that is translated to multiple IP addresses. When an outside host accesses one of the mapped IP addresses, it is untranslated to the single load balancer address. Depending on the URL requested, it redirects traffic to the correct web server.

Figure 360: Static NAT with One-to-Many for an Inside Load Balancer



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

Procedure

Step 1 Create a network object for the addresses to which you want to map the load balancer.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the addresses.

Name the network object (for example, myPublicIPs) and enter the network range 209.165.201.3-209.165.201.5.

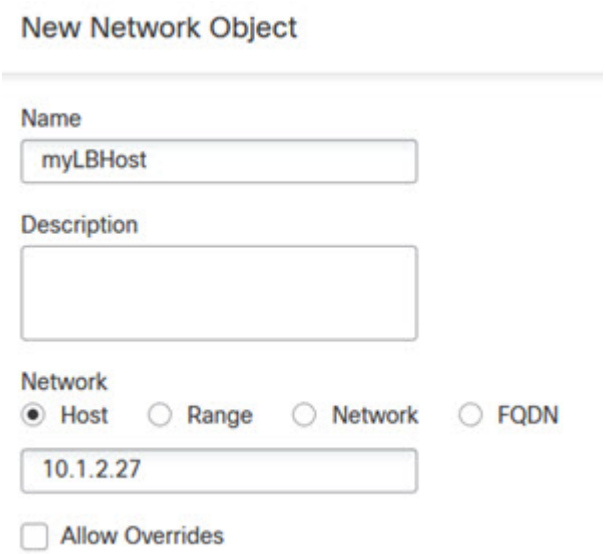


d) Click **Save**.

Step 2 Create a network object for the load balancer.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, myLBHost), enter the host address 10.1.2.27.



c) Click **Save**.

Step 3 Configure static NAT for the load balancer.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

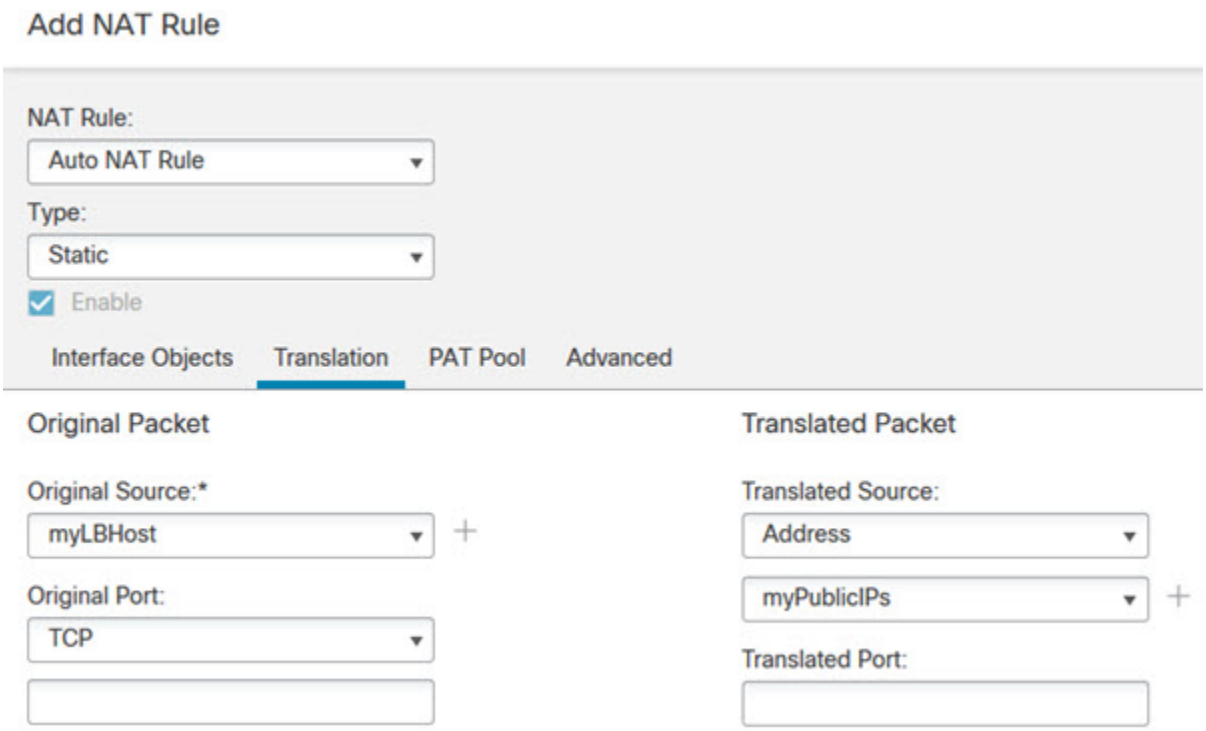
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = myLBHost network object.
- **Translated Source** > **Address** = myPublicIPs network group.



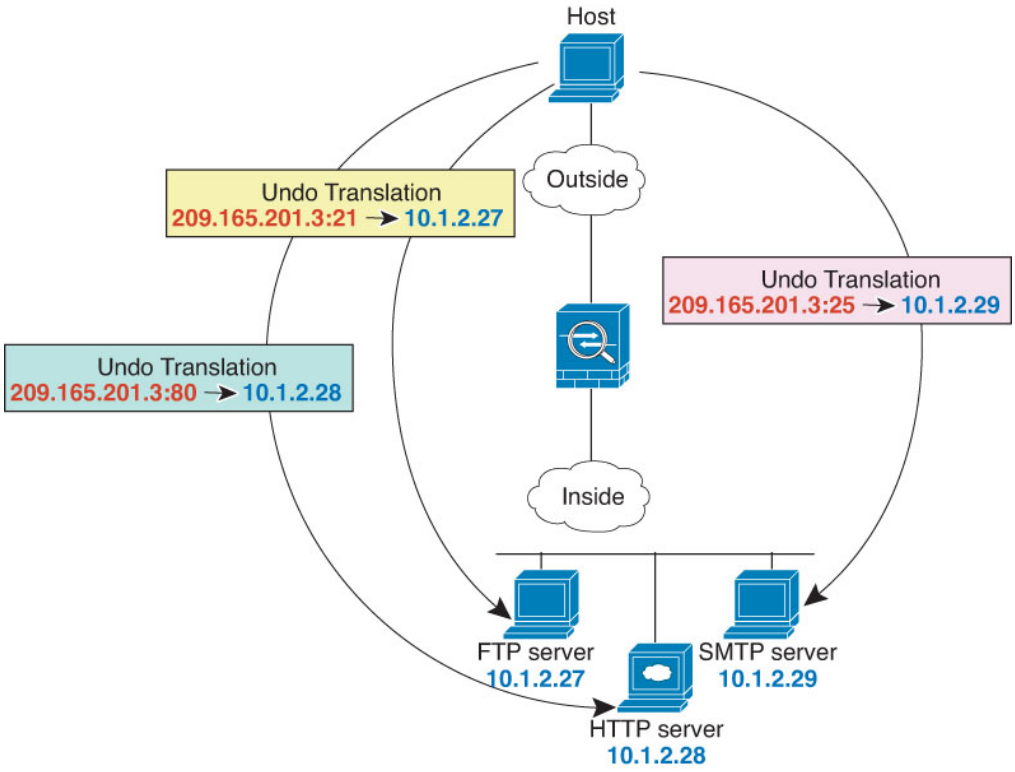
f) Click **Save**.

Step 4 Click **Save** on the NAT rule page.

Single address for FTP, HTTP, and SMTP (static auto NAT with port translation)

The following static NAT-with-port-translation example provides a single address for remote users to access FTP, HTTP, and SMTP. These servers are actually different devices on the real network, but for each server, you can specify static NAT-with-port-translation rules that use the same mapped IP address, but different ports.

Figure 361: Static NAT-with-Port-Translation



Before you begin

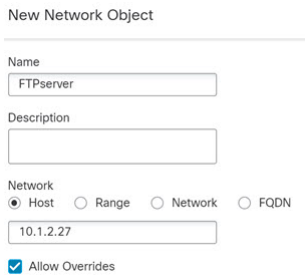
Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select and click **Add Network** > **Add GroupObjects** > **Object Management** > **Interface**.

Procedure

Step 1 Create a network object for the FTP server.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Name the network object (for example, FTPserver), and enter the real IP address for the FTP server, 10.1.2.27.

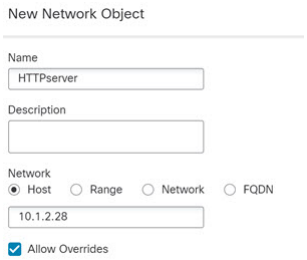


c) Click **Save**.

Step 2 Create a network object for the HTTP server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, HTTPserver), enter the host address 10.1.2.28.

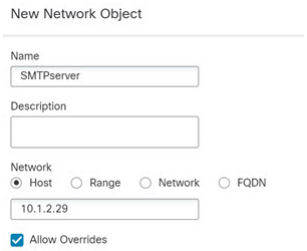


c) Click **Save**.

Step 3 Create a network object for the SMTP server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, SMTPserver), enter the host address 10.1.2.29.

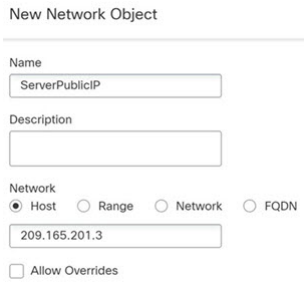


c) Click **Save**.

Step 4 Create a network object for the public IP address used for the three servers.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, ServerPublicIP) and enter the host address 209.165.201.3.



c) Click **Save**.

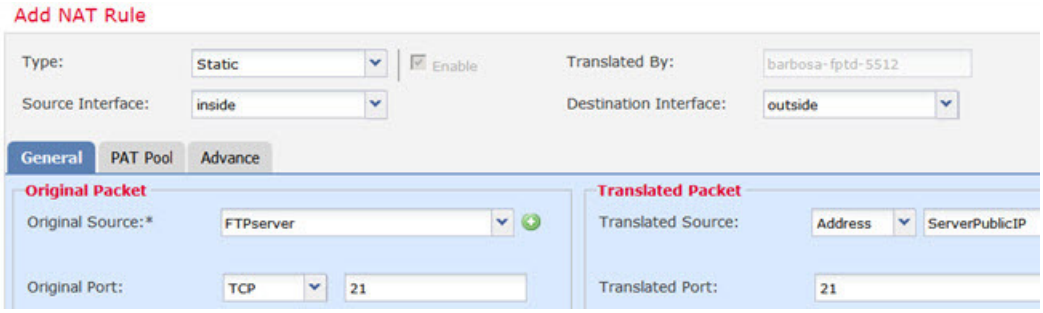
Step 5 Configure static NAT with port translation for the FTP server, mapping the FTP port to itself.

a) Select **Devices** > **Device Management** and edit the Threat Defense device on which you are defining the rule. Select **NAT**.

Ensure that you are currently in the same leaf domain that includes the device, if you are using domains.

b) Click **Add NAT Rule** > **Add Auto NAT**.
c) Configure the following properties:

- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Source** = FTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 21.
- **Translated Port** = 21.



d) Click **Save**.

Step 6 Configure static NAT with port translation for the FTP server, mapping the FTP port to itself.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:

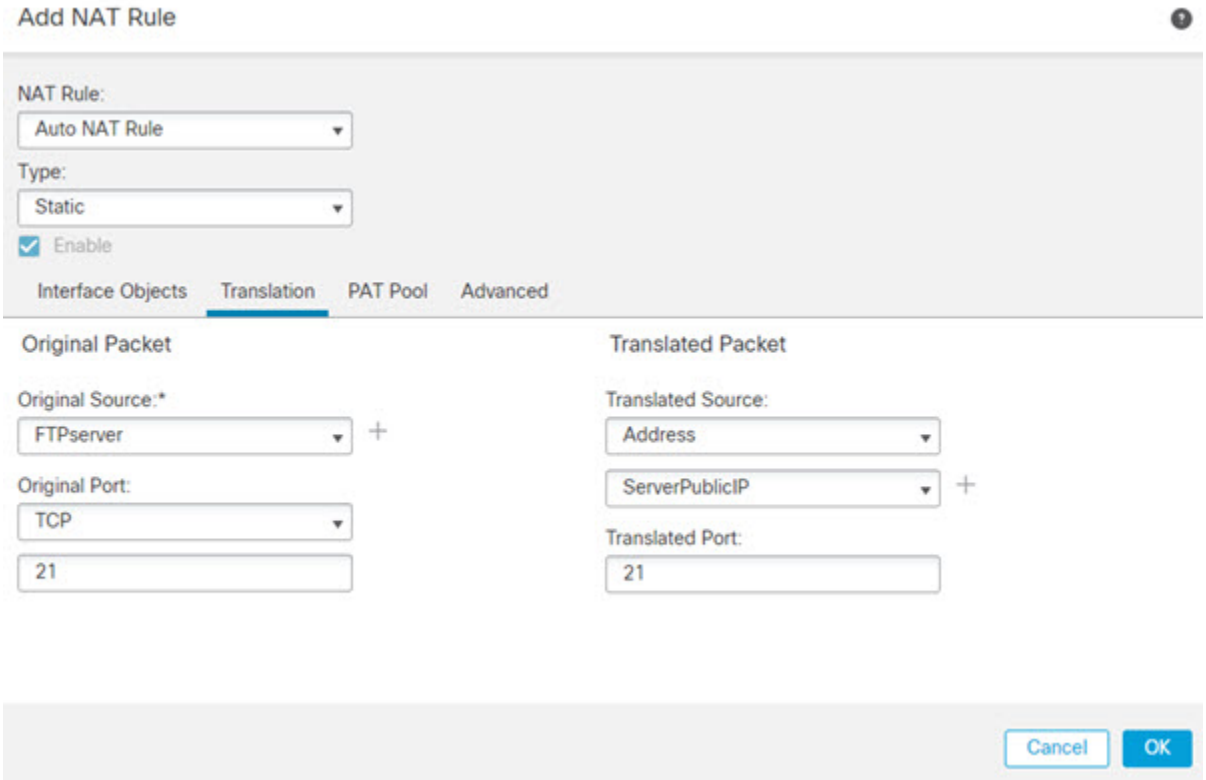
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = FTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 21.
- **Translated Port** = 21.



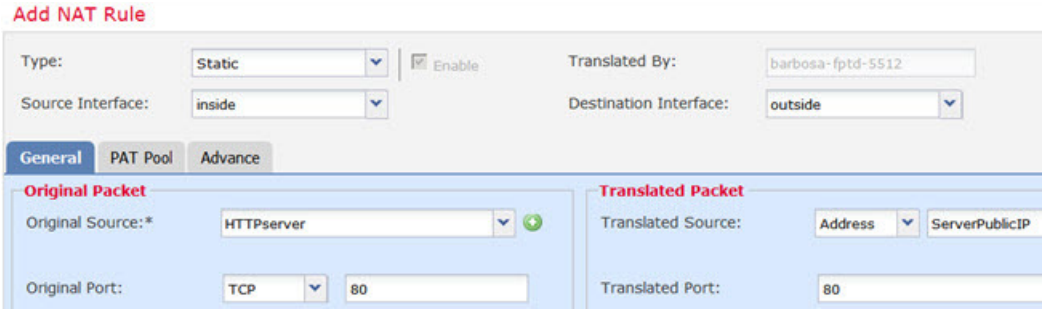
f) Click **Save**.

Step 7 Configure static NAT with port translation for the HTTP server, mapping the HTTP port to itself.

a) Click **Add NAT Rule** > **Add Auto NAT**.

b) Configure the following properties:

- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Source** = HTTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 80.
- **Translated Port** = 80.



c) Click **Save**.

Step 8 Configure static NAT with port translation for the HTTP server, mapping the HTTP port to itself.

a) Click **Add Rule**.

b) Configure the following properties:

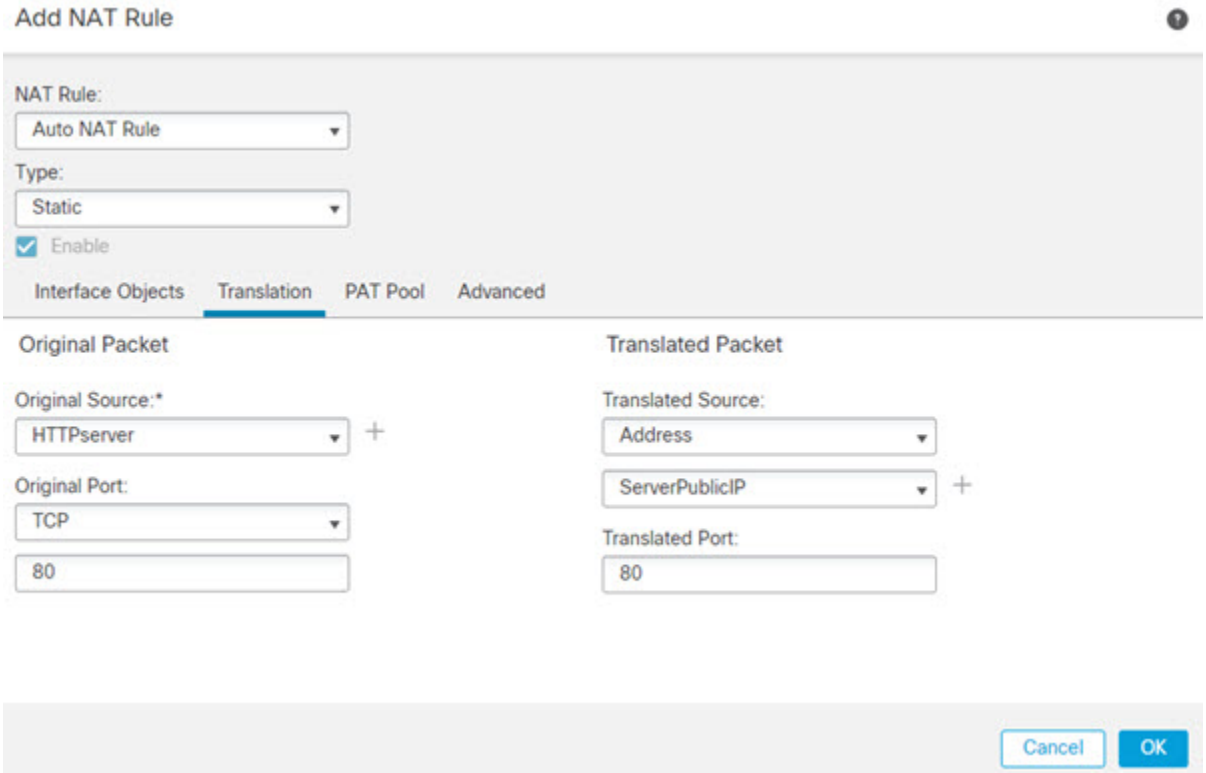
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = HTTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 80.
- **Translated Port** = 80.



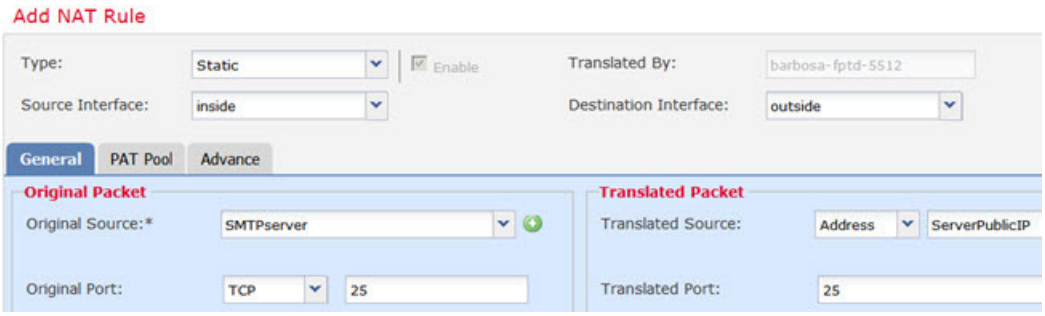
e) Click **Save**.

Step 9 Configure static NAT with port translation for the SMTP server, mapping the SMTP port to itself.

a) Click **Add NAT Rule** > **Add Auto NAT**.

b) Configure the following properties:

- **Type** = Static.
- **Source Interface** = inside.
- **Destination Interface** = outside.
- **Original Source** = SMTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 25.
- **Translated Port** = 25.



c) Click **Save**.

Step 10 Configure static NAT with port translation for the SMTP server, mapping the SMTP port to itself.

a) Click **Add Rule**.

b) Configure the following properties:

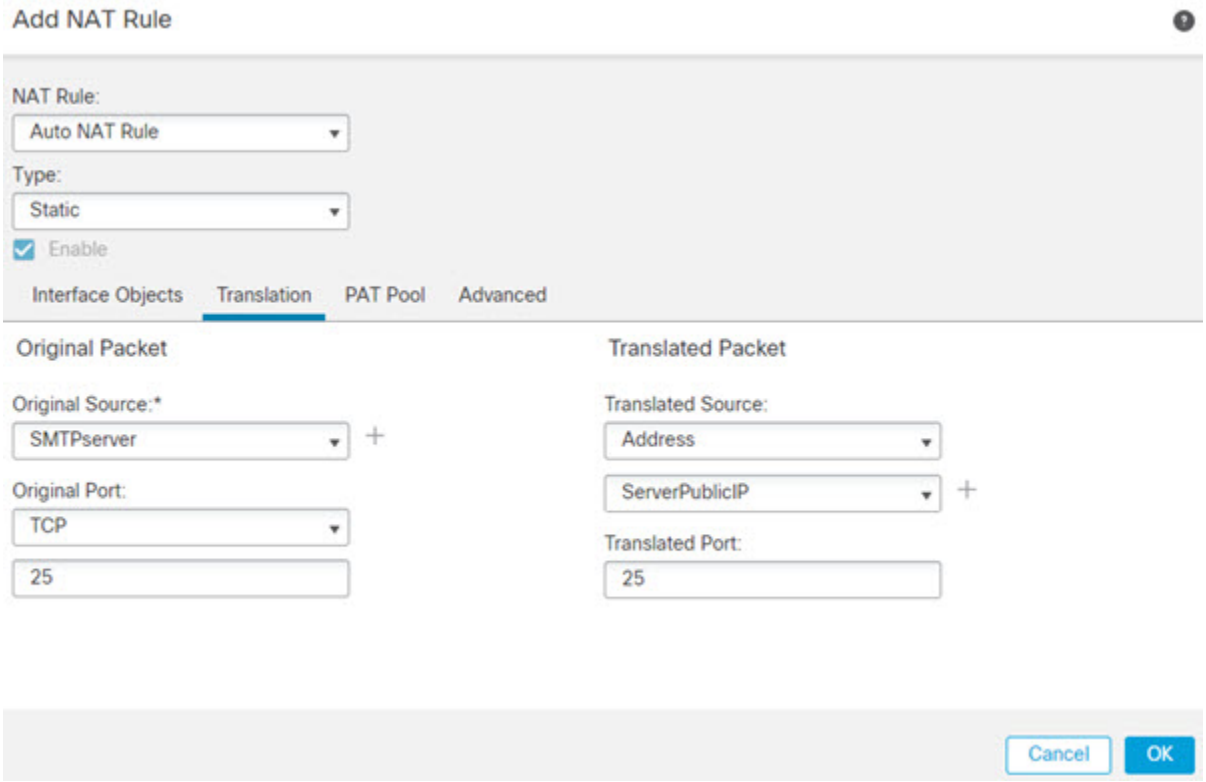
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = SMTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 25.
- **Translated Port** = 25.



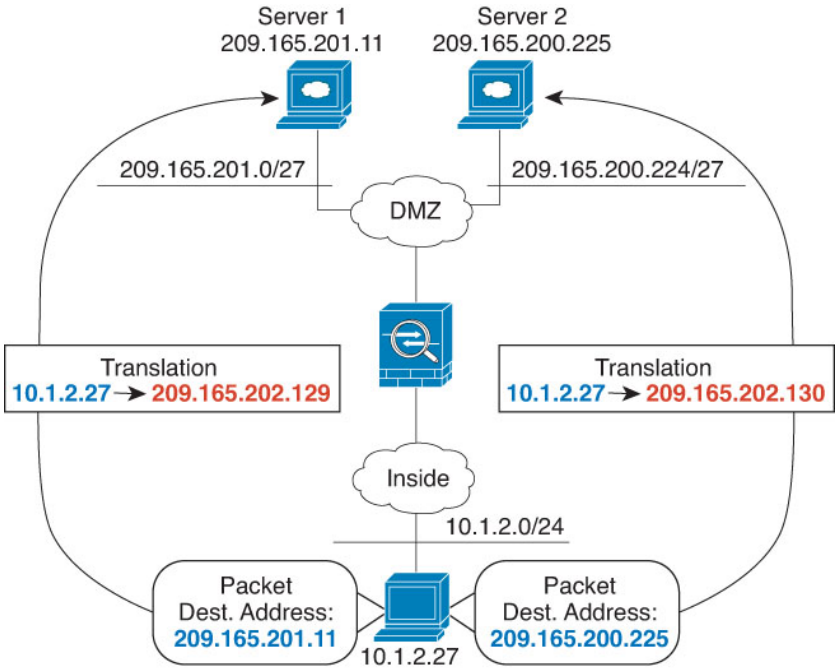
e) Click **Save**.

Step 11 Click **Save** on the NAT rule page.

Different translation depending on the destination (dynamic manual PAT)

The following figure shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129:*port*. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130:*port*.

Figure 362: Manual NAT with Different Destination Addresses



Before you begin

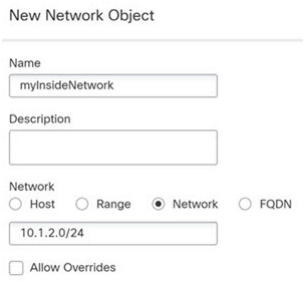
Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **dmz**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Name the network object (for example, myInsideNetwork), and enter the real network address, 10.1.2.0/24.



c) Click **Save**.

Step 2 Create a network object for the DMZ network 1.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, DMZnetwork1) and enter the network address 209.165.201.0/27 (subnet mask of 255.255.255.224).



c) Click **Save**.

Step 3 Create a network object for the PAT address for DMZ network 1.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress1) and enter the host address 209.165.202.129.



c) Click **Save**.

Step 4 Create a network object for the DMZ network 2.

a) Click **Add Network** > **Add Object**.

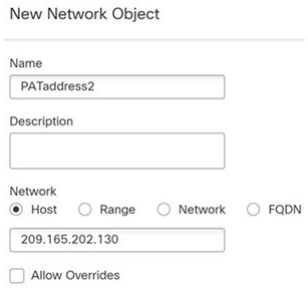
b) Name the network object (for example, DMZnetwork2) and enter the network address 209.165.200.224/27 (subnet mask of 255.255.255.224).



c) Click **Save**.

Step 5 Create a network object for the PAT address for DMZ network 2.

a) Click **Add Network** > **Add Object**.
b) Name the network object (for example, PATaddress2) and enter the host address 209.165.202.130.



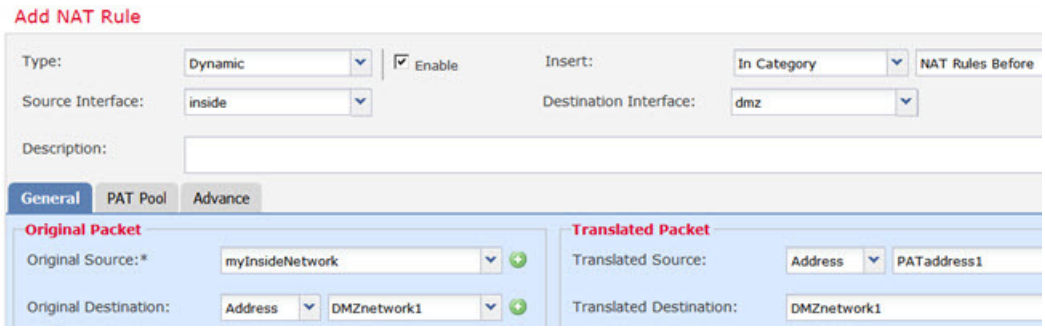
c) Click **Save**.

Step 6 Configure dynamic manual PAT for DMZ network 1.

a) Select **Devices** > **Device Management** and edit the Threat Defense device on which you are defining the rule. Select **NAT**.
b) Click **Add NAT Rule** > **Add Manual NAT**.
c) Configure the following properties:

- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress1 network object.
- **Original Destination** > **Address** = DMZnetwork1 network object.
- **Translated Destination** = DMZnetwork1 network object.

Note
Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank.



d) Click **Save**.

Step 7 Configure dynamic manual PAT for DMZ network 1.

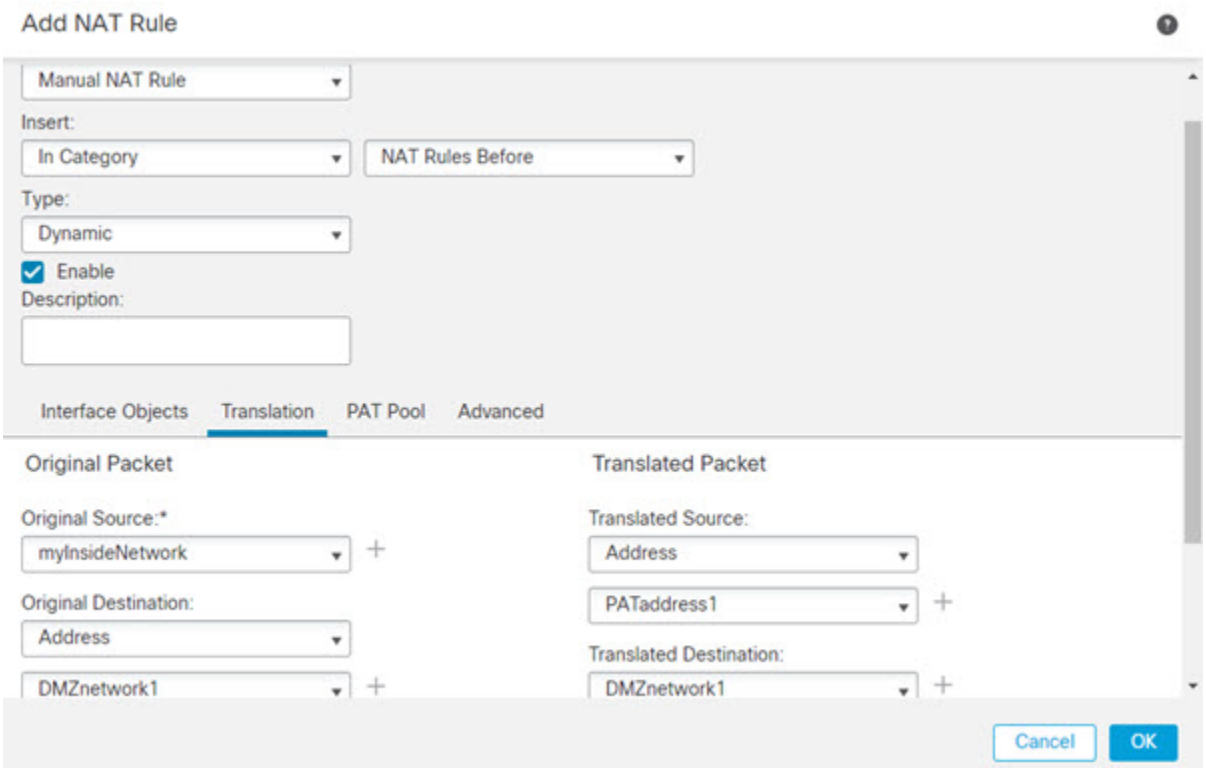
a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.
c) Configure the following properties:
   - **NAT Rule** = Manual NAT Rule.
   - **Type** = Dynamic.

d) On **Interface Objects**, configure the following:
   - **Source Interface Objects** = inside.
   - **Destination Interface Objects** = dmz.

e) On **Translation**, configure the following:
   - **Original Source** = myInsideNetwork network object.
   - **Translated Source** > **Address**= PATaddress1 network object.
   - **Original Destination** > **Address** = DMZnetwork1 network object.
   - **Translated Destination** = DMZnetwork1 network object.

   Note
   Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank.



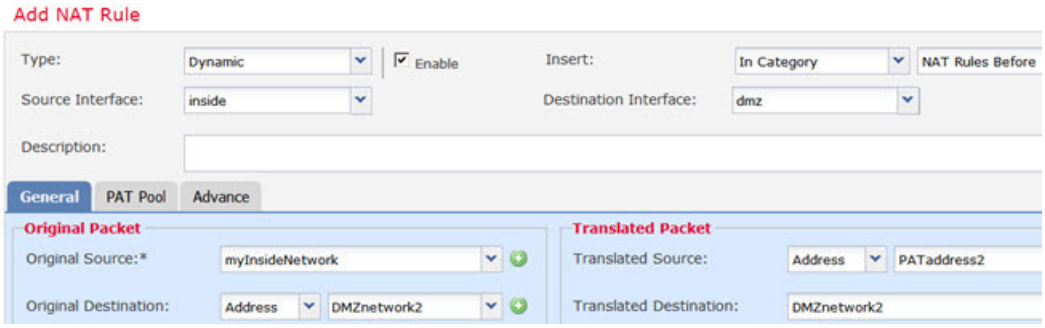
f) Click **Save**.

Step 8 Configure dynamic manual PAT for DMZ network 2.

a) Click **Add NAT Rule** > **Add Manual NAT**.

b) Configure the following properties:

- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress2 network object.
- **Original Destination** > **Address** = DMZnetwork2 network object.
- **Translated Destination** = DMZnetwork2 network object.



c) Click **Save**.

Step 9 Configure dynamic manual PAT for DMZ network 2.

a) Click **Add Rule**.

b) Configure the following properties:

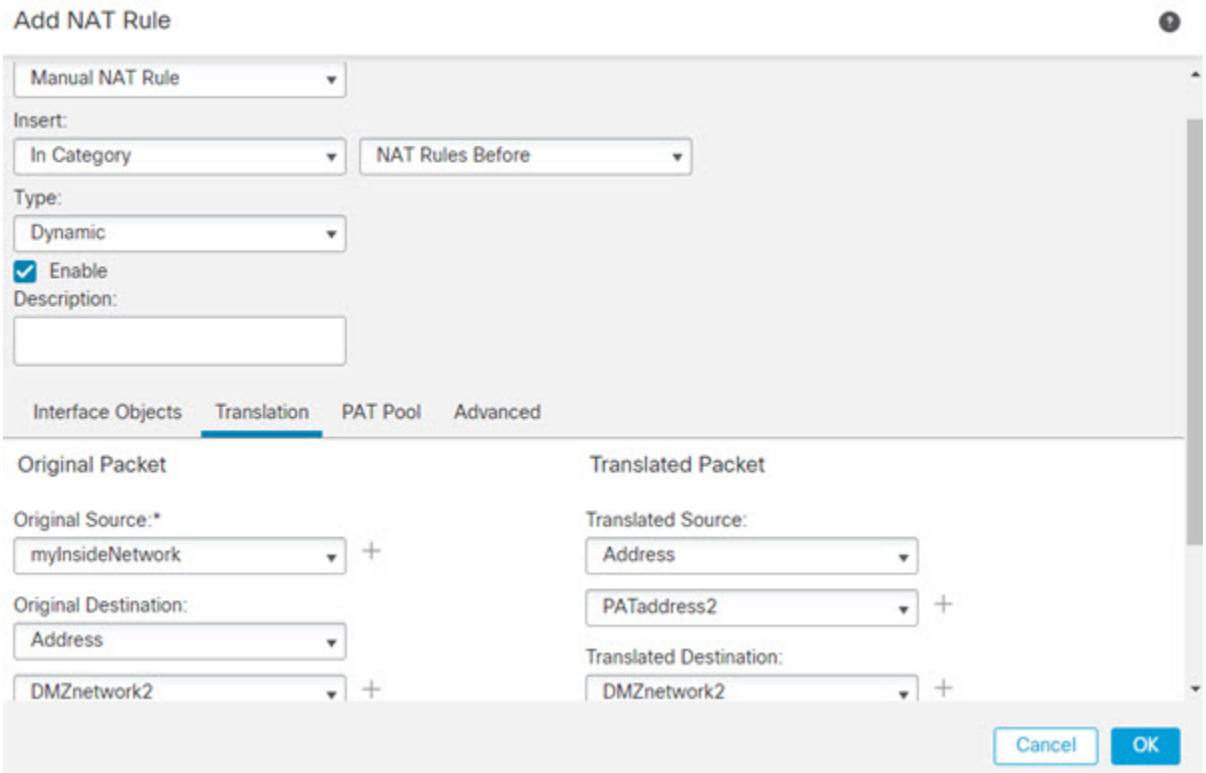
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

d) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress2 network object.
- **Original Destination** > **Address** = DMZnetwork2 network object.
- **Translated Destination** = DMZnetwork2 network object.



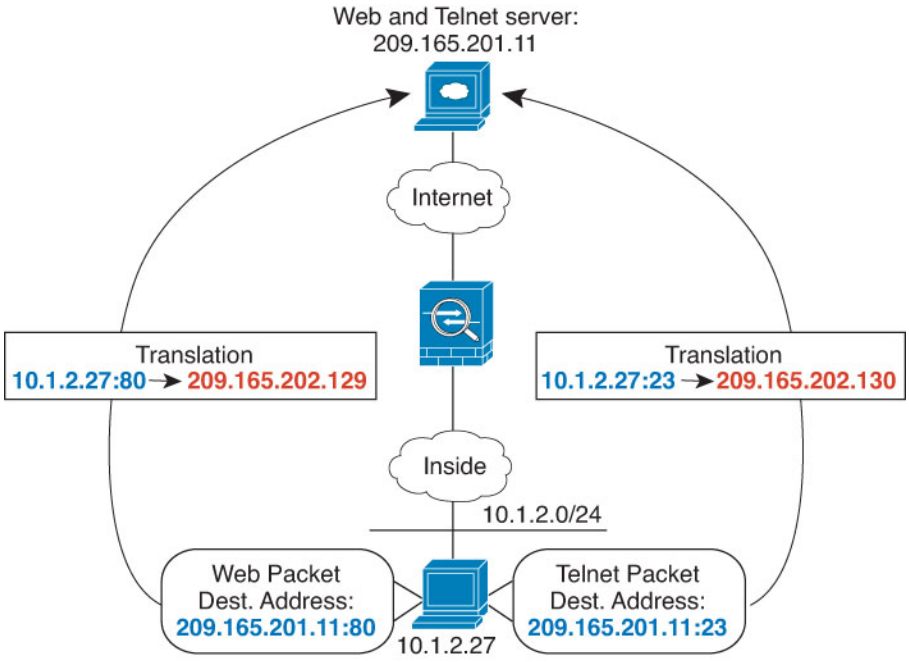
e) Click **Save**.

Step 10 Click **Save** on the NAT rule page.

Configure dynamic manual PAT with different destination ports

This illustration displays the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for Telnet services, the real address is translated to 209.165.202.129: *port*. When the host accesses the same server for web services, the real address is translated to 209.165.202.130: *port*.

Figure 363: Manual NAT with Different Destination Ports



Before you begin

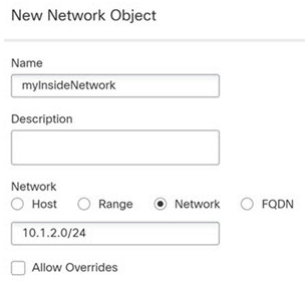
Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **dmz**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Name the network object (for example, myInsideNetwork) and enter the real network address, 10.1.2.0/24.



c) Click **Save**.

Step 2 Create a network object for the Telnet/Web server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, TelnetWebServer) and enter the host address 209.165.201.11.

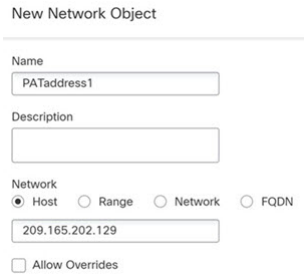


c) Click **Save**.

Step 3 Create a network object for the PAT address when using Telnet.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress1) and enter the host address 209.165.202.129.

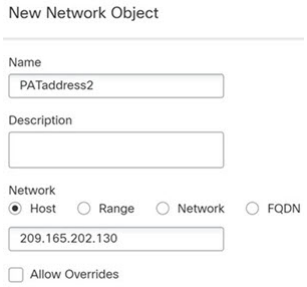


c) Click **Save**.

Step 4 Create a network object for the PAT address when using HTTP.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress2) and enter the host address 209.165.202.130.



c) Click **Save**.

Step 5 Configure dynamic manual PAT for Telnet access.

a) Select **Devices** > **Device Management** and edit the Threat Defense device on which you are defining the rule. Select **NAT**.

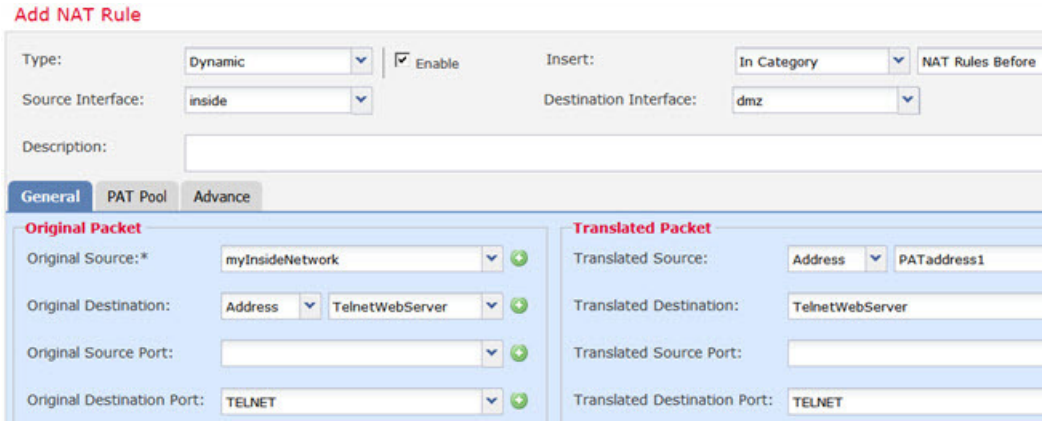
b) Click **Add NAT Rule** > **Add Manual NAT**.

c) Configure the following properties:

- **Type** = Dynamic.
- **Source Interface** = inside.

- **Destination Interface** = dmz.
- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address** = PATaddress1 network object.
- **Original Destination** > **Address** = TelnetWebServer network object.
- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = TELNET port object (system-defined).
- **Translated Destination Port** = TELNET port object (system-defined).

Note
Because you do not want to translate the destination address or port, you need to configure identity NAT for them by specifying the same address for the original and translated destination addresses, and the same port for the original and translated port.



d) Click **Save**.

Step 6 Configure dynamic manual PAT for Telnet access.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

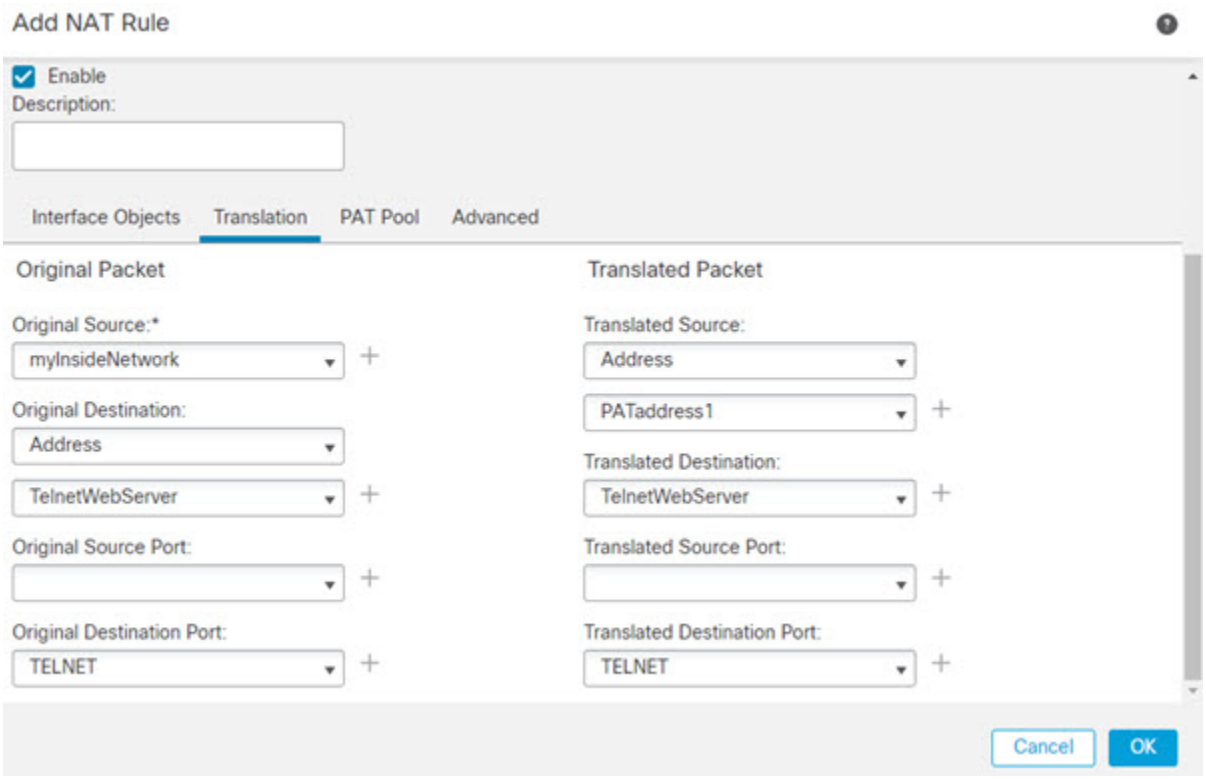
- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

e) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address** = PATaddress1 network object.

- **Original Destination** > **Address** = TelnetWebServer network object.
- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = TELNET port object (system-defined).
- **Translated Destination Port** = TELNET port object (system-defined).

Note
Because you do not want to translate the destination address or port, you need to configure identity NAT for them by specifying the same address for the original and translated destination addresses, and the same port for the original and translated port.



f) Click **Save**.

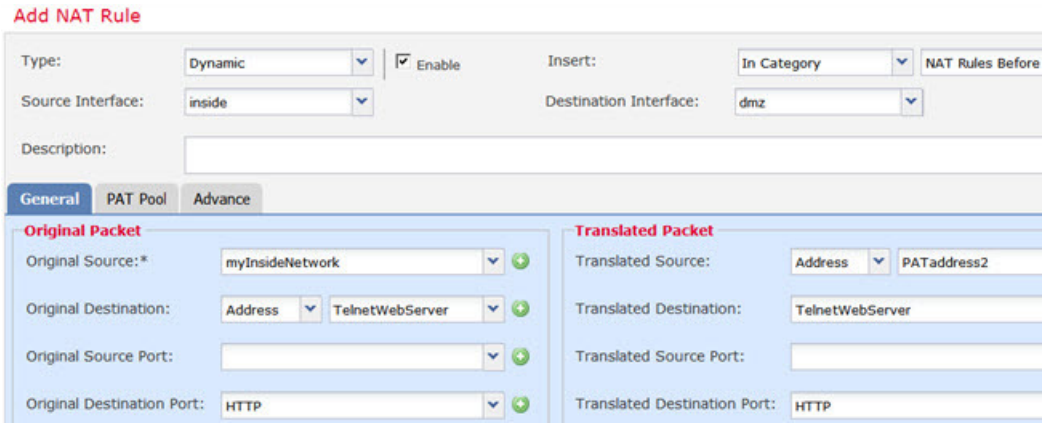
Step 7 Configure dynamic manual PAT for web access.

a) Click **Add NAT Rule** > **Add Manual NAT**.

b) Configure the following properties:

- **Type** = Dynamic.
- **Source Interface** = inside.
- **Destination Interface** = dmz.
- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address** = PATaddress2 network object.
- **Original Destination** > **Address** = TelnetWebServer network object.

- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = HTTP port object (system-defined).
- **Translated Destination Port** = HTTP port object (system-defined).



c) Click **Save**.

Step 8 Configure dynamic manual PAT for web access.

a) Click **Add Rule**.

b) Configure the following properties:

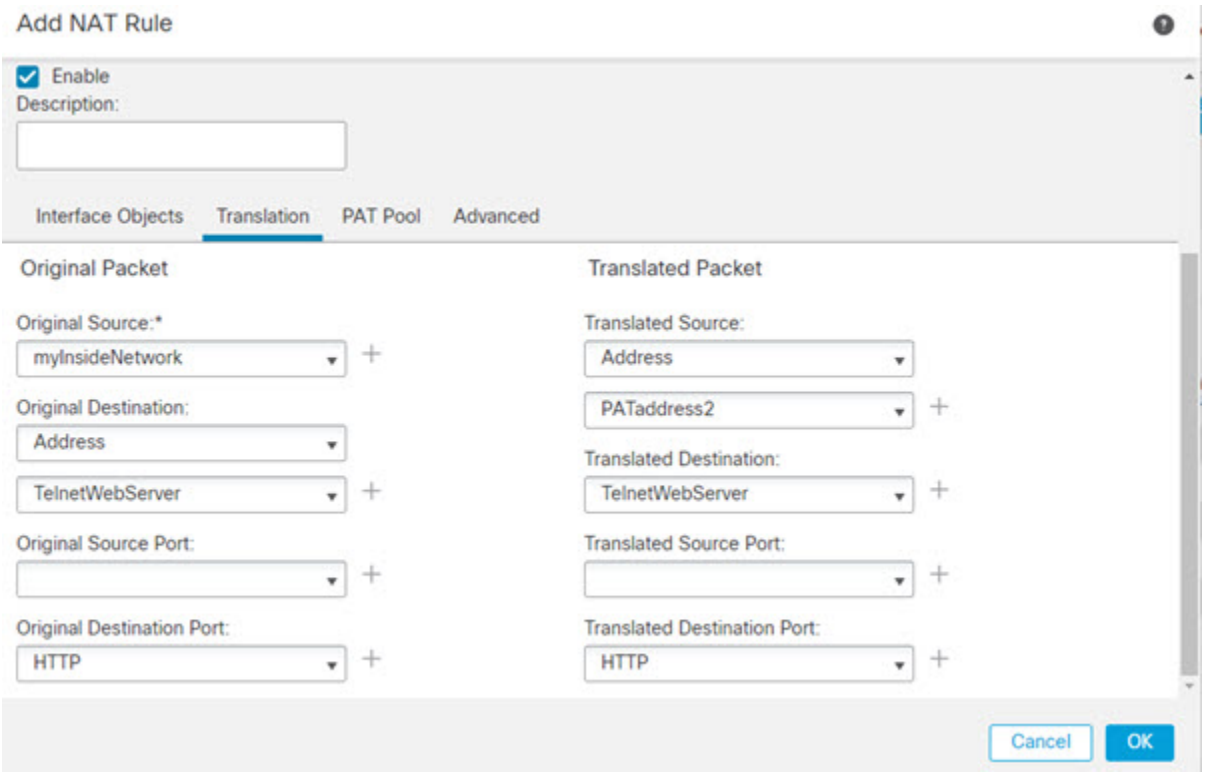
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

d) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address** = PATaddress2 network object.
- **Original Destination** > **Address** = TelnetWebServer network object.
- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = HTTP port object (system-defined).
- **Translated Destination Port** = HTTP port object (system-defined).



e) Click **Save**.

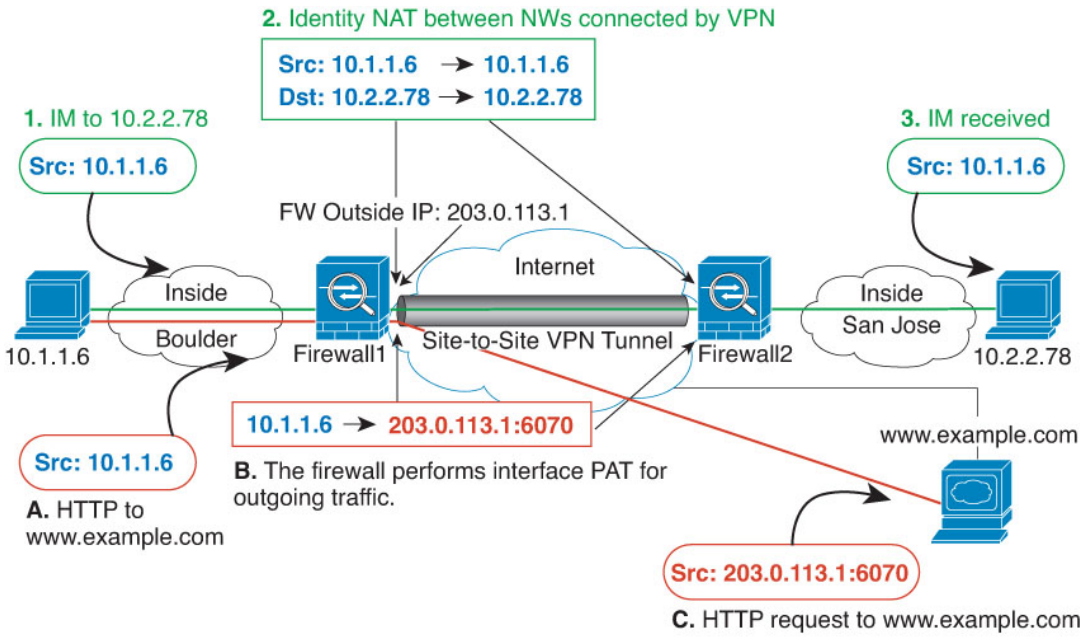
Step 9 Click **Save** on the NAT rule page.

NAT and site-to-site VPN

When you create a policy-based site-to-site VPN using the management center VPN wizard (**Device > Site To Site**), you can select the NAT Exempt option to create the rules automatically. You can view the NAT exemptions for a device in the NAT policy page (**Device > NAT > NAT Exemptions**). If you do not want to configure NAT Exempt in the VPN wizard, you can use the following procedure for NAT exemption.

The following figure shows a site-to-site tunnel connecting the Boulder and San Jose offices. For traffic that you want to go to the Internet (for example from 10.1.1.6 in Boulder to www.example.com), you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. However, for traffic that you want to go over the VPN tunnel (for example from 10.1.1.6 in Boulder to 10.2.2.78 in San Jose), you do not want to perform NAT; you need to exempt that traffic by creating an identity NAT rule. Identity NAT simply translates an address to the same address.

Figure 364: Interface PAT and Identity NAT for Site-to-Site VPN



This example explains the configuration for Firewall1 (Boulder).

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the devices in the VPN. In this example, we will assume the interface objects are security zones named **inside-boulder** and **outside-boulder** for the Firewall1 (Boulder) interfaces. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

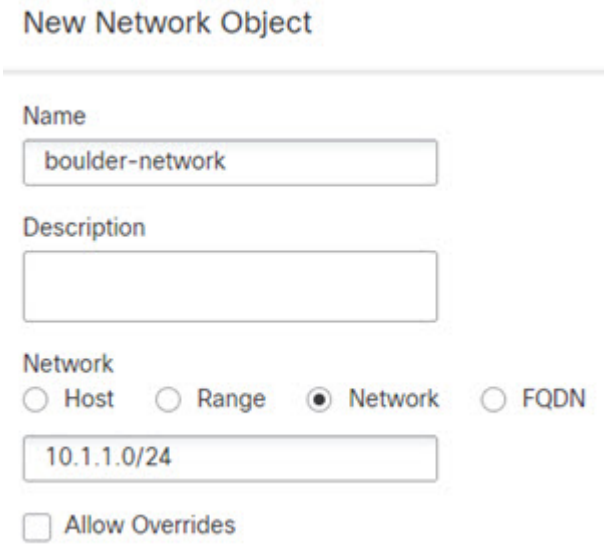
Procedure

Step 1 Create the objects to define the various networks.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Identify the Boulder inside network.

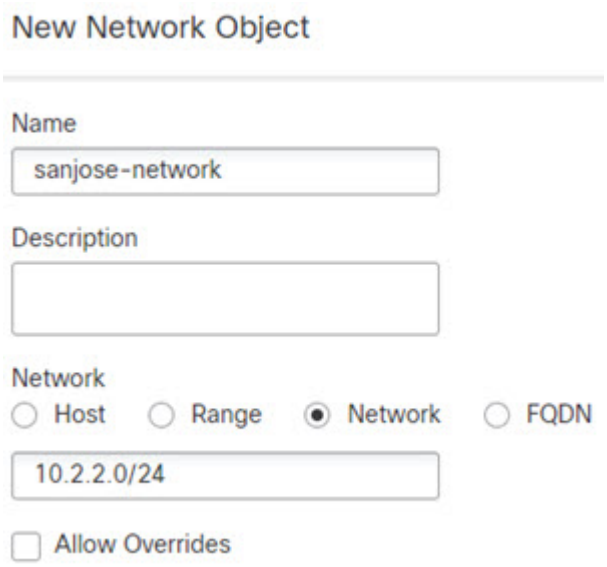
Name the network object (for example, boulder-network) and enter the network address, 10.1.1.0/24.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the inside San Jose network.

Name the network object (for example, sanjose-network) and enter the network address 10.2.2.0/24.



e) Click **Save**.

Step 2 Configure manual identity NAT for the Boulder network when going over the VPN to San Jose on Firewall1 (Boulder).

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.

- **Type** = Static.

d) On **Interface Objects**, configure the following:

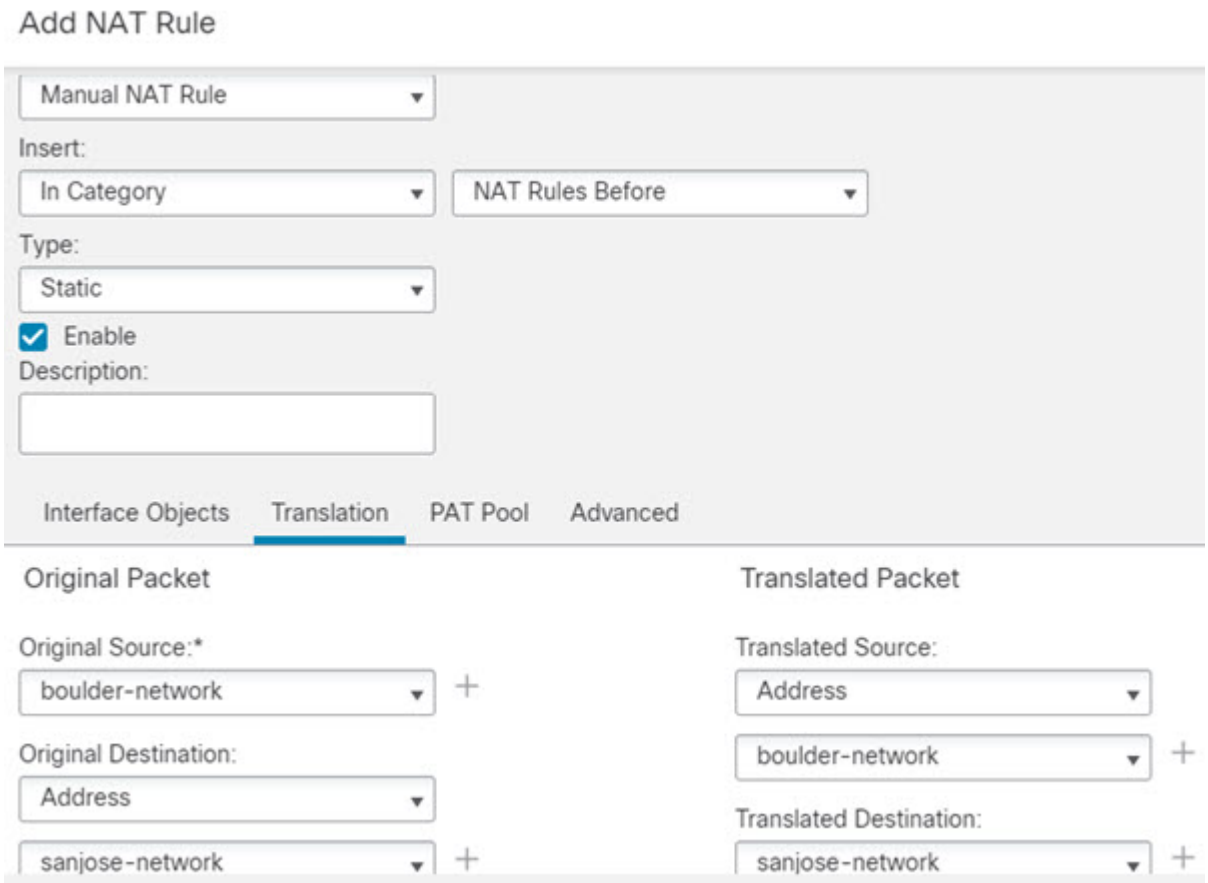
- **Source Interface Objects** = inside-boulder.
- **Destination Interface Objects** = outside-boulder.

e) On **Translation**, configure the following:

- **Original Source** = boulder-network object.
- **Translated Source** > **Address** = boulder-network object.
- **Original Destination** > **Address** = sanjose-network object.
- **Translated Destination** = sanjose-network object.

Note
Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank. This rule configures identity NAT for both source and destination.

f) On **Advanced**, select **Do not proxy ARP on Destination interface**.



g) Click **Save**.

Step 3 Configure manual dynamic interface PAT when going to the Internet for the inside Boulder network on Firewall1 (Boulder).

a) Click **Add Rule**.

b) Configure the following properties:

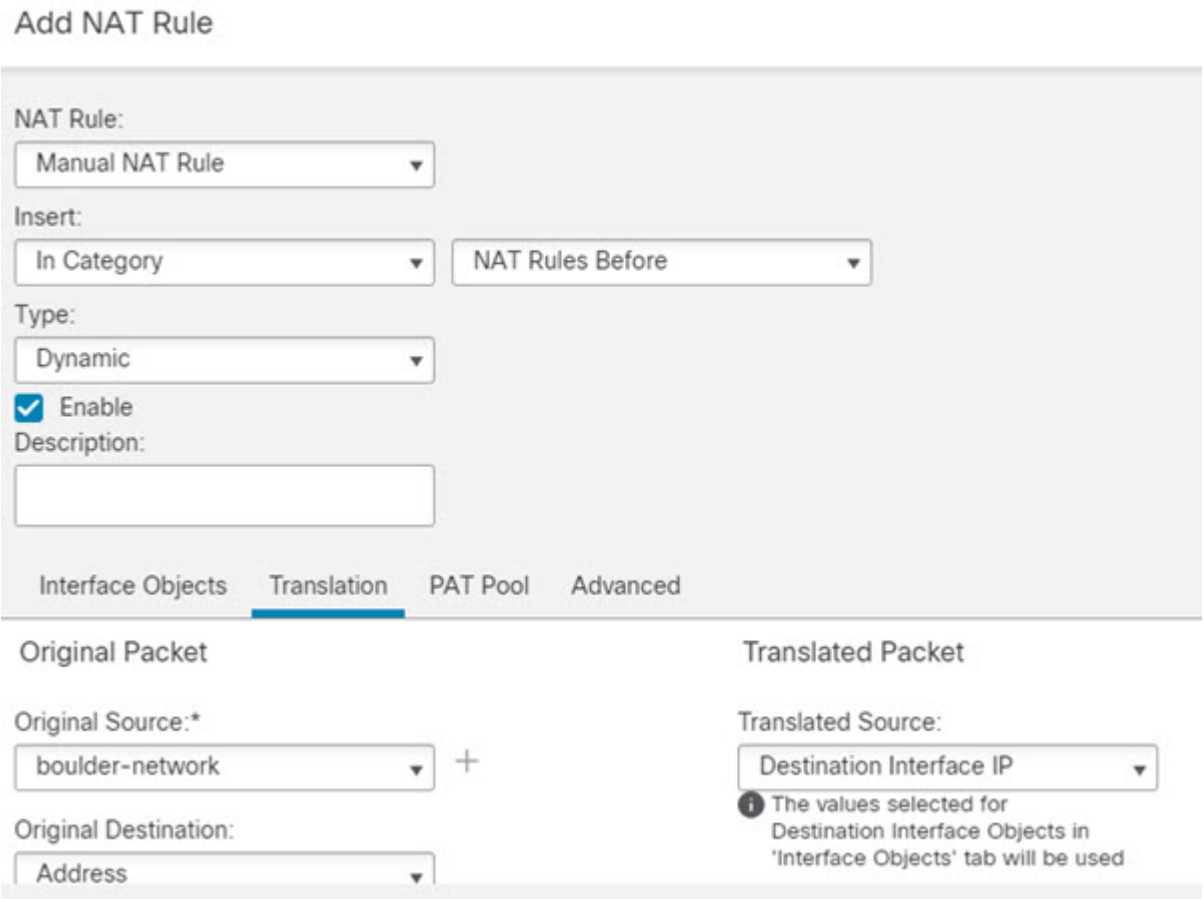
- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.
- **Insert Rule** = any position after the first rule. Because this rule will apply to any destination address, the rule that uses sanjose-network as the destination must come before this rule, or the sanjose-network rule will never be matched. The default is to place new manual NAT rules at the end of the "NAT Rules Before Auto NAT" section.

c) On **Interface Objects** , configure the following:

- **Source Interface Objects** = inside-boulder.
- **Destination Interface Objects** = outside-boulder.

d) On **Translation** , configure the following:

- **Original Source** = boulder-network object.
- **Translated Source** = **Destination Interface IP**. This option configures interface PAT using the interface contained in the destination interface object.
- **Original Destination** > **Address** = any (leave blank).
- **Translated Destination** = any (leave blank).



e) Click **Save**.

Step 4 If you are also managing Firewall2 (San Jose), you can configure similar rules for that device.

- The manual identity NAT rule would be for sanjose-network when the destination is boulder-network. Create new interface objects for the Firewall2 inside and outside networks.
- The manual dynamic interface PAT rule would be for sanjose-network when the destination is "any."

DNS query and response rewriting using NAT

DNS rewrite is the process of modifying DNS replies so that the IP address returned in the DNS response is replaced with an address that matches the NAT configuration. It is used when a translated address must be returned to the requester instead of the original address in the DNS record. This function is also known as DNS doctoring.

This feature rewrites the address in DNS queries and replies that match a NAT rule (for example, the A record for IPv4, the AAAA record for IPv6, or the PTR record for reverse DNS queries). For DNS replies traversing from a mapped interface to any other interface, the record is rewritten from the mapped value to the real value. Inversely, for DNS replies traversing from any interface to a mapped interface, the record is rewritten from the real value to the mapped value. This feature works with NAT44,NAT 66, NAT46, and NAT64.

The main circumstances when you would need to configure DNS rewrite on a NAT rule are:

- The rule is NAT64 or NAT46, and the DNS server is on the outside network. You need DNS rewrite to convert between DNS A records (for IPv4) and AAAA records (for IPv6).
- The DNS server is on the outside, clients are on the inside, and some of the fully-qualified domain names that the clients use resolve to other inside hosts.
- The DNS server is on the inside and responds with private IP addresses, clients are on the outside, and the clients access fully-qualified domain names that point to servers that are hosted on the inside.

DNS rewrite has the following limitations:

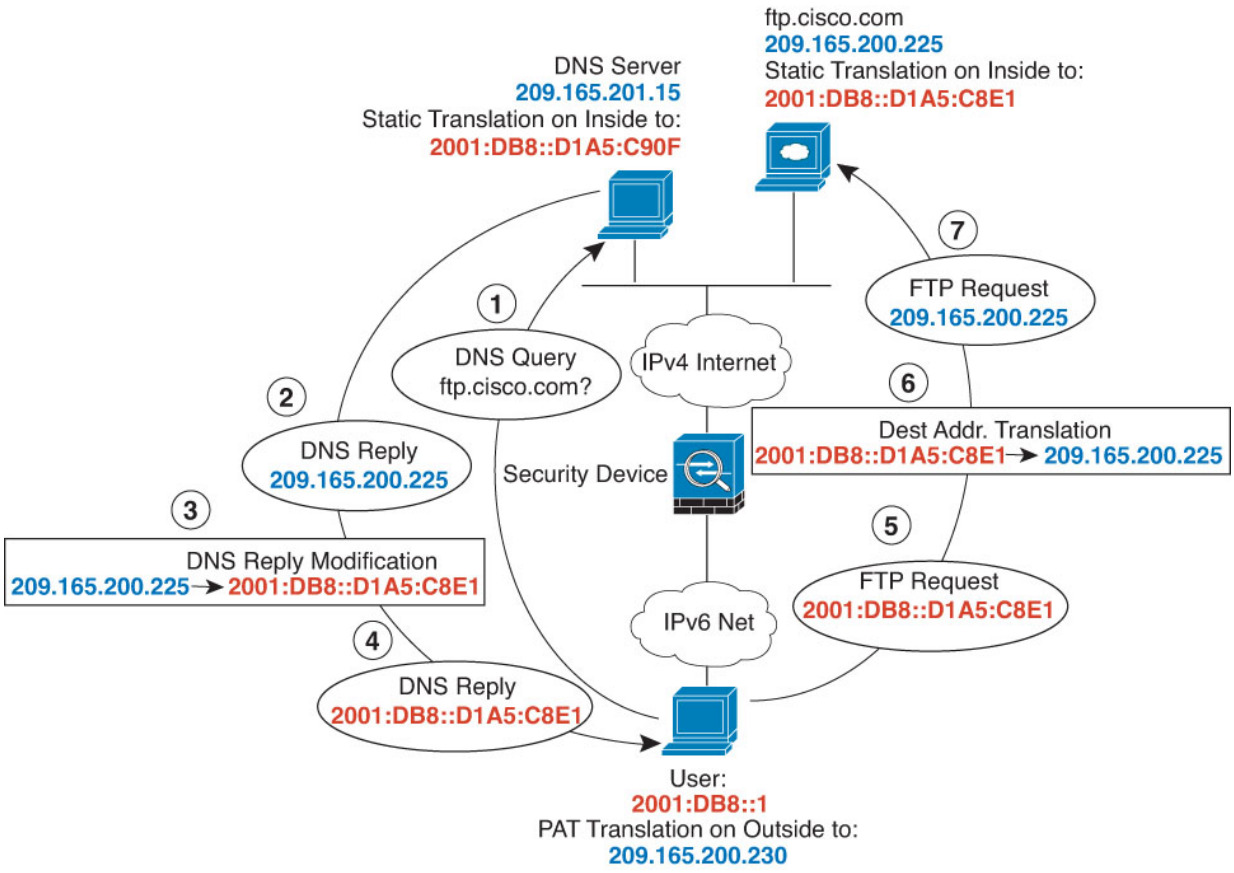
- DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A or AAAA record, and the PAT rule to use is ambiguous.
- If you configure a manual NAT rule, you cannot configure DNS modification if you specify the destination address as well as the source address. These kinds of rules can potentially have a different translation for a single address when going to A vs. B. Therefore, the can not accurately match the IP address inside the DNS reply to the correct twice NAT rule; the DNS reply does not contain information about which source/destination address combination was in the packet that prompted the DNS request.
- You must enable DNS application inspection with DNS NAT rewrite enabled for NAT rules to rewrite DNS queries and responses. By default, DNS inspection with DNS NAT rewrite enabled is globally applied, so you probably do not need to change the inspection configuration.
- DNS rewrite is actually done on the xlate entry, not the NAT rule. Thus, if there is no xlate for a dynamic rule, rewrite cannot be done correctly. The same problem does not occur for static NAT.
- DNS rewrite does not rewrite DNS Dynamic Update messages (opcode 5).

These topics provide examples of DNS rewrite in NAT rules.

DNS64 reply modification

The following figure shows an FTP server and DNS server on the outside IPv4 network. The system has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225.

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1, where D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

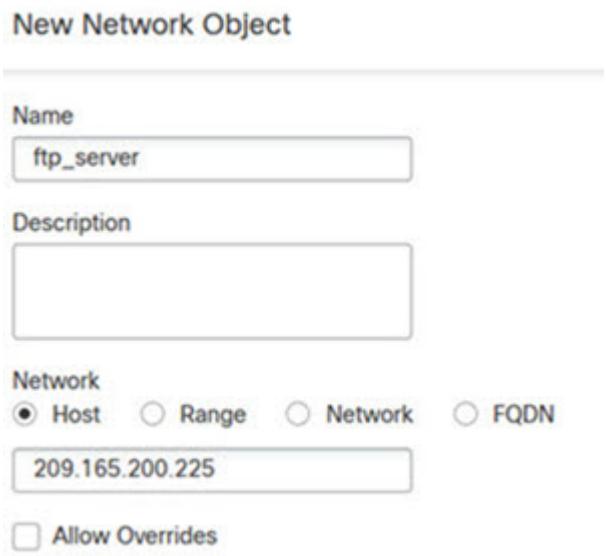
Procedure

Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the real FTP server address.

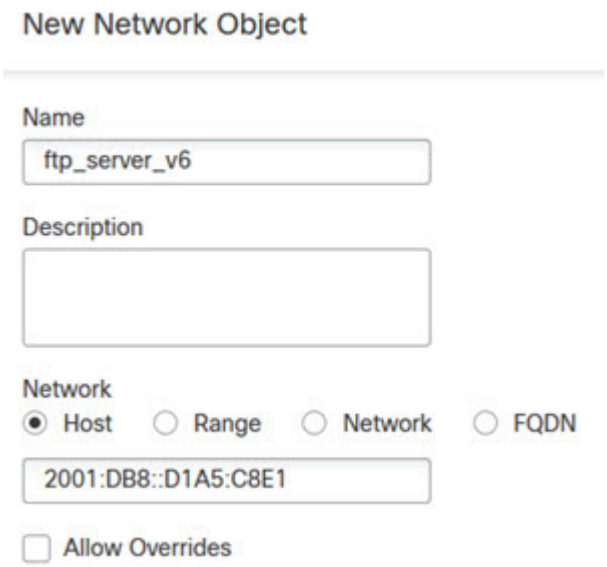
Name the network object (for example, ftp_server) and enter the host address, 209.165.200.225.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the FTP server's translated IPv6 address.

Name the network object (for example, ftp_server_v6) and enter the host address, 2001:DB8::D1A5:C8E1.

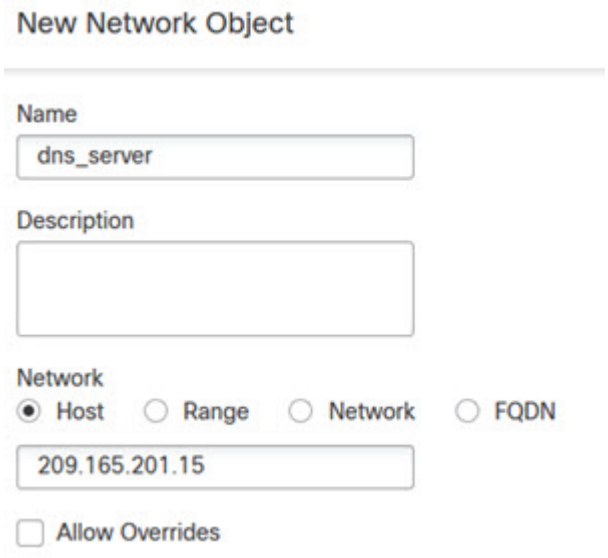


e) Click **Save**.

Step 2 Create the network objects for the DNS server.

a) Click **Add Network** > **Add Object** and define the DNS server's real address.

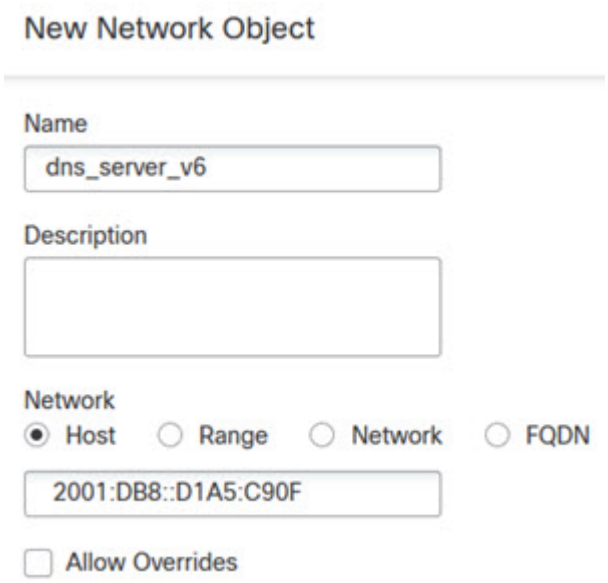
Name the network object (for example, dns_server) and enter the host address, 209.165.201.15.



b) Click **Save**.

c) Click **Add Network** > **Add Object** and define the DNS server's translated IPv6 address.

Name the network object (for example, dns_server_v6) and enter the host address, 2001:DB8::D1A5:C90F (where D1A5:C90F is the IPv6 equivalent of 209.165.201.15).

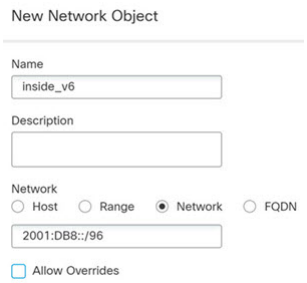


d) Click **Save**.

Step 3 Create the network objects for the inside network.

a) Click **Add Network** > **Add Object** and define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:DB8::/96.

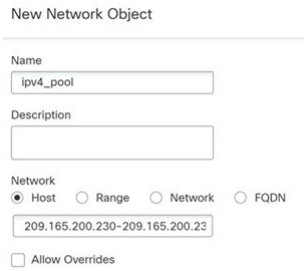


b) Click **Save**.

Step 4 Create the network objects for the PAT pool.

a) Click **Add Network** > **Add Object** and define the IPv4 PAT pool for the inside IPv6 network.

Name the network object (for example, ipv4_pool) and enter the range 209.165.200.230-209.165.200.235.



b) Click **Save**.

Step 5 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

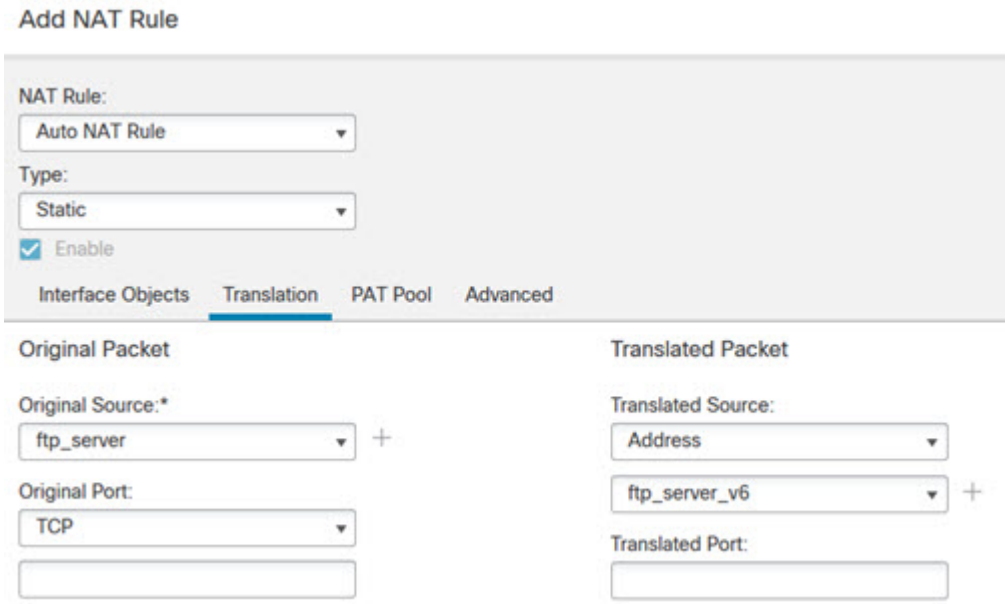
- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_v6 network object.



f) On **Advanced**, select the following options:

- **Translate DNS replies that match this rule**.
- **Net to Net Mapping**, because this is a one-to-one NAT46 translation.

g) Click **OK**.

Step 6 Configure the static NAT rule for the DNS server.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

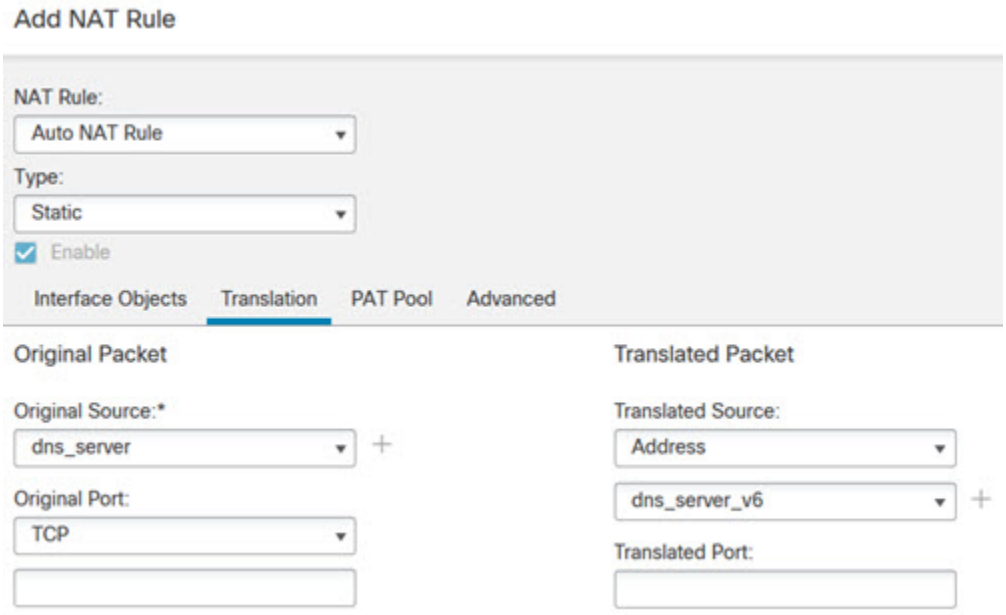
c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = dns_server network object.
- **Translated Source** > **Address** = dns_server_v6 network object.

e) On **Advanced**, select **Net to Net Mapping**, because this is a one-to-one NAT46 translation.



f) Click **OK**.

Step 7 Configure the dynamic NAT with a PAT pool rule for the inside IPv6 network.

a) Click **Add Rule**.

b) Configure the following properties:

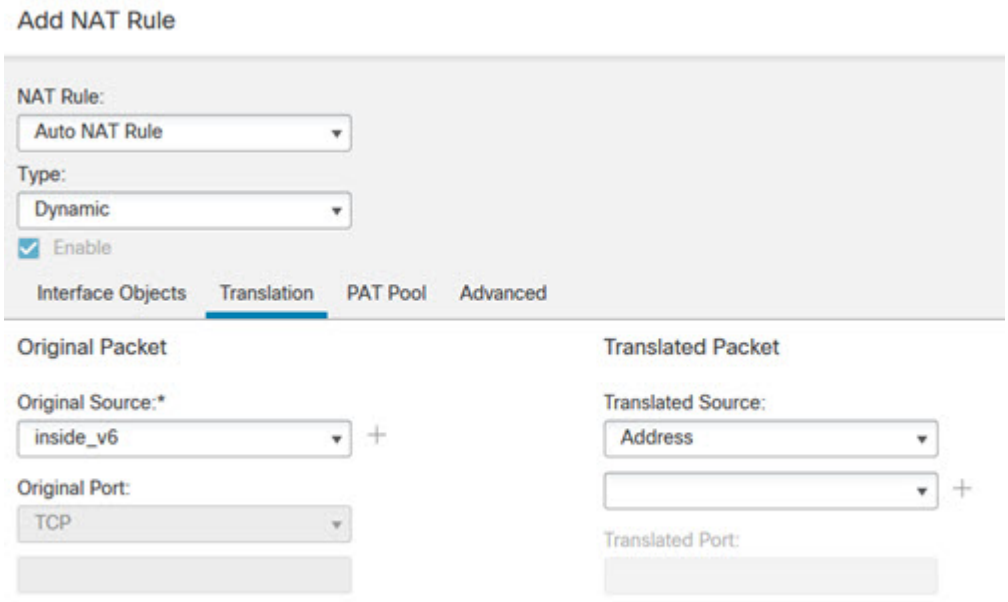
- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

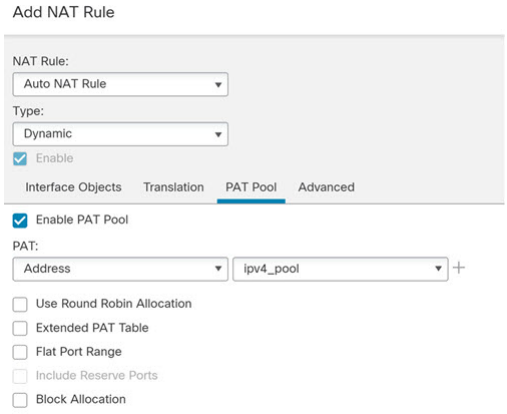
d) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** > **Address** = leave this field empty.



e) On **PAT Pool**, configure the following:

- **Enable PAT Pool** = select this option.
- **Translated Source** > **Address** = ipv4_pool network object.



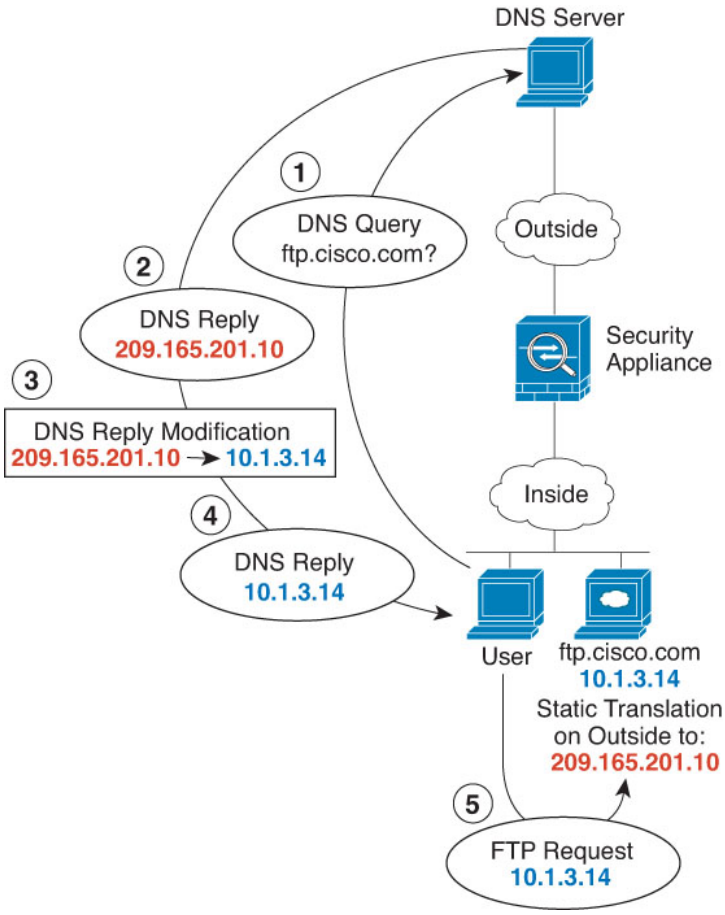
f) Click **OK**.

Create DNS reply modification with DNS server on outside

The following figure shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure NAT to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network.

In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The system refers to the static rule for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

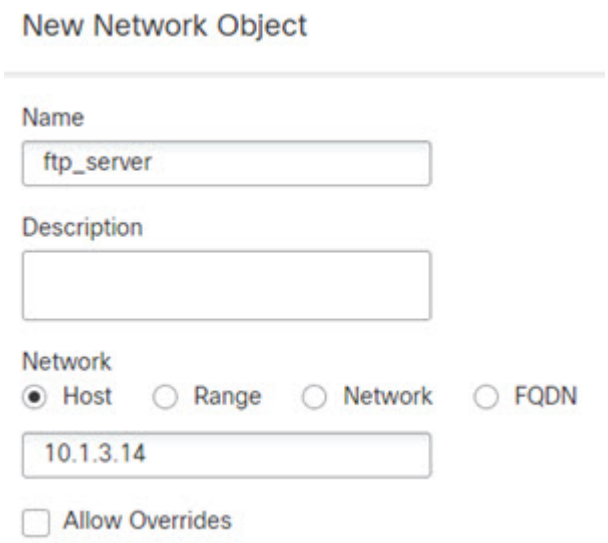
Procedure

Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the real FTP server address.

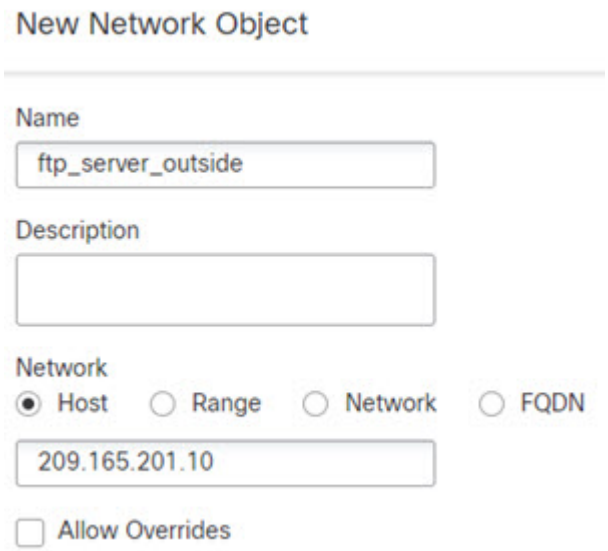
Name the network object (for example, ftp_server) and enter the host address, 10.1.3.14.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the FTP server's translated address.

Name the network object (for example, ftp_server_outside) and enter the host address, 209.165.201.10.



e) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

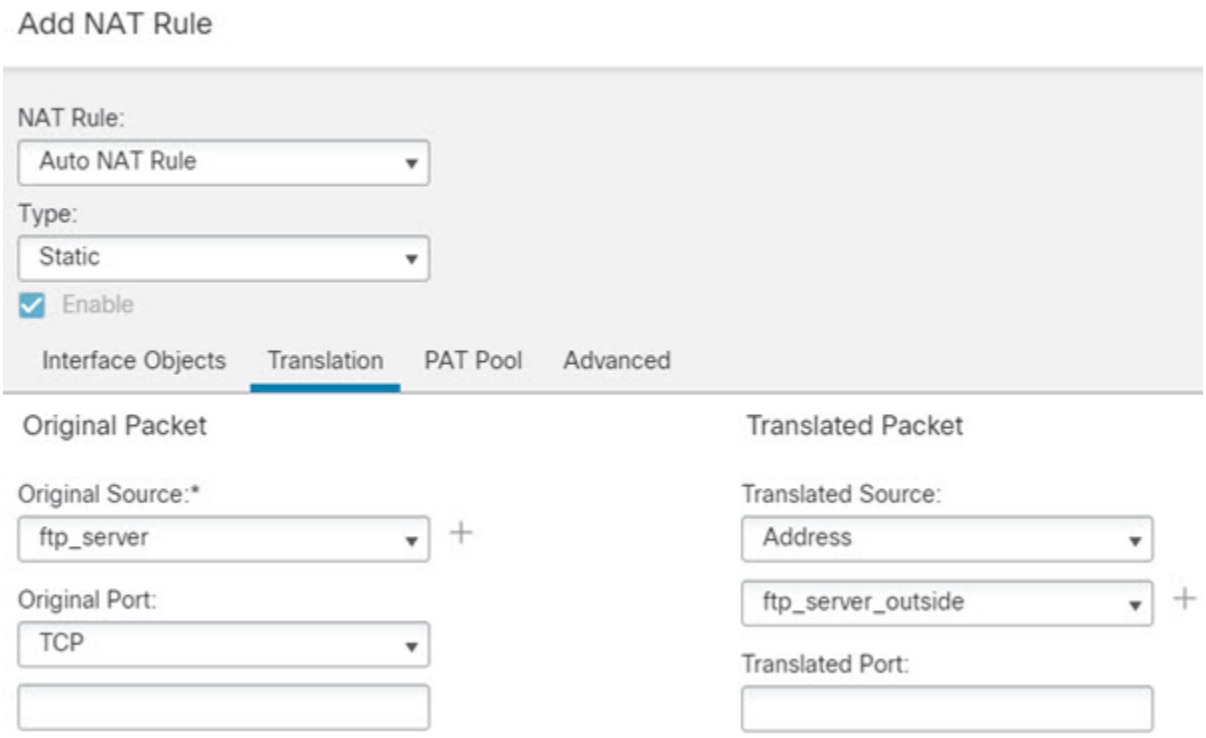
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_outside network object.

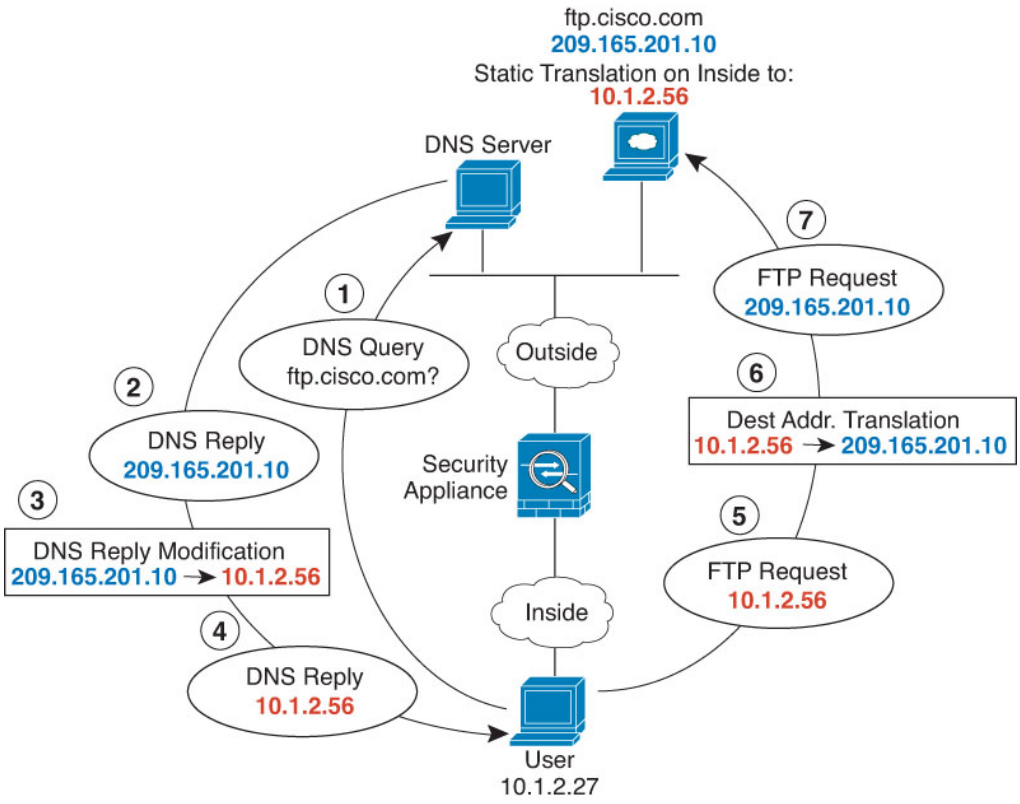
f) On **Advanced**, select **Translate DNS replies that match this rule**.



g) Click **OK**.

Configure DNS reply modification with DNS server on host network

The following figure shows an FTP server and DNS server on the outside. The system has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.



Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management** > **Interface**.

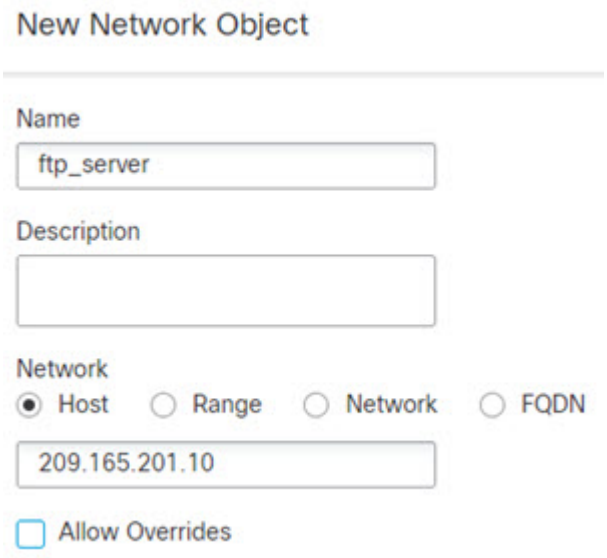
Procedure

Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management** > **Network** and click **Add Network** > **Add Object**.

b) Define the real FTP server address.

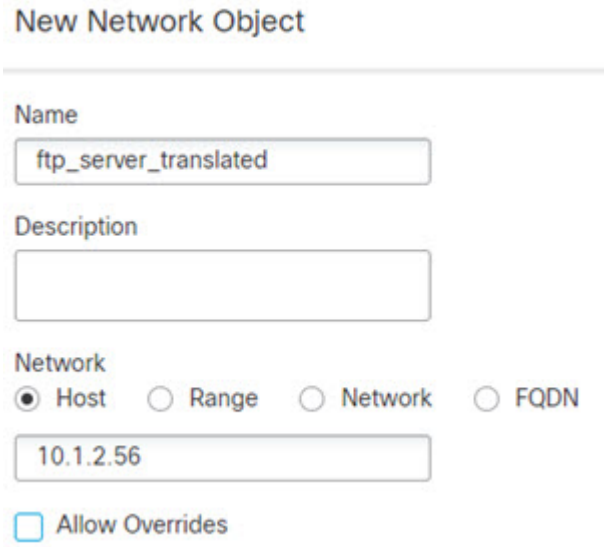
Name the network object (for example, ftp_server) and enter the host address, 209.165.201.10.



c) Click **Save**.

d) Click **Add Network** > **Add Object** and define the FTP server's translated address.

Name the network object (for example, ftp_server_translated) and enter the host address, 10.1.2.56.



e) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

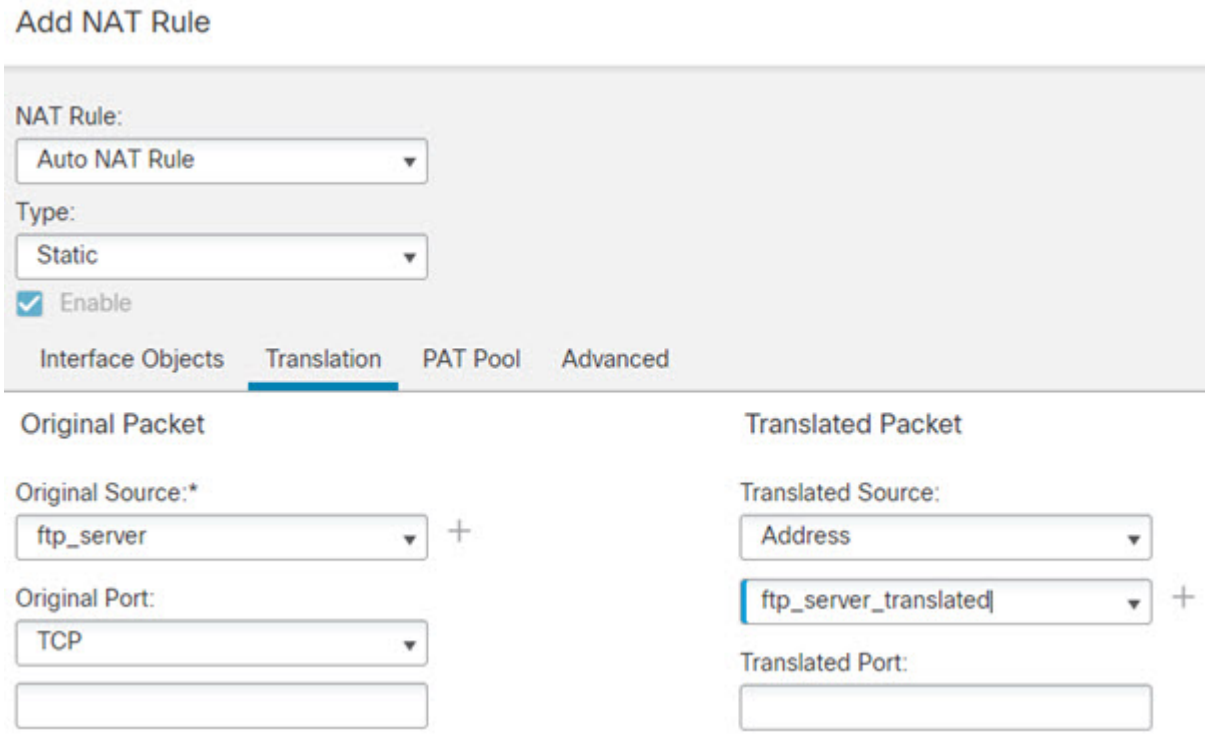
d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

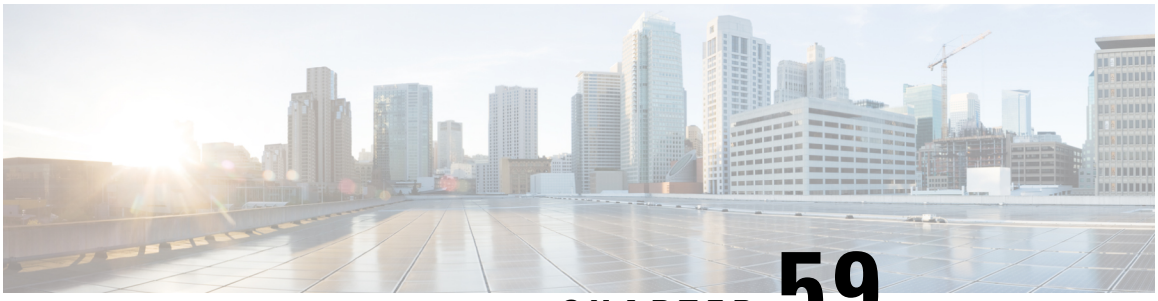
e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_translated network object.

f) On **Advanced**, select **Translate DNS replies that match this rule**.



g) Click **OK**.



CHAPTER 59

Quality of Service

Quality of Service is a network management feature that:

- enables the policing of network traffic on Firewall Threat Defense devices.
- assists administrators in managing network traffic,
- provides guidelines for rate limiting.

- QoS, on page 1693
- QoS policies, on page 1694
- Requirements and prerequisites for QoS, on page 1694
- Guidelines and limitations for QoS policies, on page 1695
- Rate limiting with QoS policies, on page 1695

QoS

QoS is a traffic management mechanism that

- rate limits (polices) network traffic that is allowed or trusted by access control
- does not rate limit traffic that was fastpathed and
- is supported only on the routed interfaces of Firewall Threat Defense devices.

QoS interface limitations

QoS is supported only on the routed interfaces of Firewall Threat Defense devices. However, it does not support site-to-site VPN and VTI interfaces.

Logging Rate-Limited Connections

You cannot configure logging for QoS. A connection can be rate limited without being logged. Conversely, a logged connection may not be rate-limited. To view QoS details, you must separately log connection ends in your Cloud-Delivered Firewall Management Center database.

Connection events for rate-limited connections show how much traffic was dropped and the specific QoS configurations that applied. You can view this information in event views (workflows), dashboards, and reports.

QoS policies

A QoS policy is a configuration mechanism that

- governs rate limiting on managed devices
- can target multiple devices with each device having one deployed QoS policy at a time, and
- contains up to 63 QoS rules that handle network traffic in specified order.

QoS policy characteristics

QoS policies have these key characteristics:

-
- The system matches traffic to QoS rules in the order specified. The system rate limits traffic according to the first rule where all rule conditions match the traffic. Traffic not matching any of the rules is not rate limited.
- QoS rules can be limited by source or destination (routed) interfaces. The system enforces rate limiting *independently* on *each* of those interfaces and does not support an aggregate rate limit for a set of interfaces.
-
- You can independently rate limit download and upload traffic based on the connection initiator.

Note A device can have no more than 255 rules, including QoS rules. Exceeding this number causes a deployment warning necessitating a reduction in the number of rules..

Note QoS operates independently and is not subordinate to a main access control configuration. However, the access control and QoS policies deployed to the same device share identity configurations. For more infomration, see Associate prefilter, decryption, and identity policies with an access control, on page 2016.

Requirements and prerequisites for QoS

This reference provides the requirements and prerequisites for implementing QoS, including supported models, domains, and user roles.

Model support

Firewall Threat Defense

Supported domains

Any

User roles

Admin

Access Admin

Network Admin

Guidelines and limitations for QoS policies

QoS rules and service policy rules (in the access control policy) are translated into firewall commands. There can be only one policy map per interface as the system combines QoS and service policy configurations automatically.

However, you can also create policy maps manually using FlexConfig. For lower-end firewalls, such as the Secure Firewall 3100, you can create up to 64 policy maps using FlexConfig. For more powerful firewalls, such as the Firepower 4100, you can create up to 128 policy maps. Ensure not to exceed this limit to prevent deployment failures. You can check the device configuration using the CLI or the deployment transcript to see how rules are translated into commands.

Rate limiting with QoS policies

To perform policy-based rate limiting, configure and deploy QoS policies to managed devices. Each QoS policy can target multiple devices; each device can have one deployed QoS policy at a time.

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Procedure

Step 1 Choose **Devices** > **QoS**.

Step 2 Click **New Policy** to create a new QoS policy and, optionally assign target devices; see Create a QoS policy, on page 1696.

You can also **Copy** ( ) or **Edit** ( ) an existing policy.

Step 3 Configure QoS rules; see Configuring QoS rules, on page 1698 and QoS rule conditions, on page 1700.

The QoS policy editor lists each rule in evaluation order and displays a summary of the rule conditions and rate limiting configurations. Use the right-click menu to manage rules, including moving, enabling, and disabling them.

This feature is helpful in larger deployments. You can **Filter by Device** to display only the rules that affect a specific device or group of devices. You can also search for rules, and the system matches text entered in the **Search Rules** field to rule names and condition values, including objects and object groups.

Note
Properly creating and ordering rules is a complex but essential to effective deployment. Careful planning is necessary to prevent rules from preempting each other, requiring additional licenses, or containing invalid

configurations. Icons in the interface represent comments, warnings, and errors. If issues exist, click **Show Warnings** to display a list. For more information, see Best practices for access control rules, on page 2058.

Step 4 Click **Policy Assignments** to identify which managed devices the policy targets; see Set target devices for a QoS policy , on page 1696.

Verify your choices, if you identified target devices during policy creation.

Step 5 Save the QoS policy.

Step 6 Since this feature must allow some packets to pass, configure your system to examine those packets.

Step 7 Deploy configuration changes.

The QoS policy is configured and deployed to the target devices, enabling policy-based rate limiting for network traffic.

Create a QoS policy

Create a QoS policy to establish the foundation for implementing rate limiting and traffic management rules on network devices.

The new QoS policy you create will not perform any rate limiting unless rules are added.

Procedure

Step 1 Choose **Devices** > **QoS**.

Step 2 Click **New Policy**.

Step 3 Enter a **Name** and, optionally, a **Description**.

Step 4 (Optional) Choose the **Available Devices** where you want to deploy the policy, then click **Add to Policy**, or drag and drop to the **Selected Devices**. To narrow the devices that appear, type a search string in the **Search** field.

Devices must be assigned before policy deployment.

Step 5 Click **Save**.

A new QoS policy is created and ready for configuration with rate limiting rules.

What to do next

- Configure and deploy the QoS policy; see Rate limiting with QoS policies, on page 1695.

Set target devices for a QoS policy

Before you begin

Ensure you are logged into the QoS policy editor and have the necessary access rights to configure system policies.

A QoS policy can target multiple devices, but each device can have one deployed QoS policy at a time.

Procedure

Step 1 In the QoS policy editor, click **Policy Assignments**.

Step 2 Build your target list:

- Add—Choose one or more **Available Devices**, then either click **Add to Policy** or drag and drop into the list of **Selected Devices**.
- Delete—Click **Delete** ( ) next to a single device, or choose multiple devices, right-click, then choose **Delete Selected**.
- Search—Enter a search string in the search field. Click **Clear** ( ) to clear the search.

Step 3 Click **OK** to save policy assignments.

Step 4 Click **Save** to save the policy.

The QoS policy is successfully configured with assigned target devices and is ready for deployment.

What to do next

- Deploy configuration changes.

Filter QoS rules by device

You can filter the list of QoS policy rules based on the device to which the rules apply. A rule's interface constraints determine if the rule affects a device.

Tip Hover your pointer over a rule criterion to see its value. If the criterion represents an object with device-specific overrides, the system displays the override value when you filter the rules list by that device. If the criterion represents an object with domain-specific overrides, the system displays the override value when you filter the rules list by devices in that domain.

Procedure

Step 1 In the policy editor, click **Rules**, then click **Filter by Device**.
A list of targeted devices and device groups appears.

Step 2 Check one or more check boxes to display only the rules that apply to those devices or groups. Or, check **All** to reset and display all of the rules.

Step 3 Click **OK**.

A message appears next to the button showing the filter used. To remove the filter, click the X in the message. Alternatively, select **All** in the Filter by device dialog box.

The QoS rules list displays only the rules that apply to the selected devices or device groups. It simplifies the process of reviewing and managing device-specific QoS policies.

Configuring QoS rules

QoS rules enable you to control and rate limit network traffic on specific interfaces by configuring rule properties, conditions, and traffic limits.

When you create or edit a rule, use the upper portion of the rule editor to configure general rule properties. Use the lower portion of the rule editor to configure rule conditions and comments.

Procedure

Step 1 On Rules of the QoS policy editor:

- Add Rule—Click **Add Rule**.
- Edit Rule—Click **Edit** ( ).

Step 2 Enter a **Name**.

Step 3 Configure rule components:

- Enabled—Specify whether the rule is **Enabled**.
- Apply QoS On—Select the interfaces you want to rate limit, either **Interfaces in Destination Interface Objects** or **Interfaces in Source Interface Objects**. Your choice must correspond with a populated interface constraint (not **any**).
- Traffic Limit Per Interface—Enter a **Download Limit** and an **Upload Limit** in Mbits/sec. The default value of **Unlimited** prevent matching traffic from being rate limited in that direction.
- Conditions—Click the corresponding condition you want to add. You must configure a source or destination interface condition, corresponding to your choice for **Apply QoS On**.
- Comments—Click **Comments**. To add a comment click **New Comment**, enter a comment, and click **OK**. You can edit or delete this comment until you save the rule.

For detailed information on rule components, see QoS rule components, on page 1699.

Step 4 Save the rule.

Step 5 In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste.

Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.

Step 6 Click **Save** to save the policy.

The QoS rule you created or modified is saved to the policy and configured with your set traffic limits and conditions.

What to do next

- Deploy configuration changes.

Related Topics

Best practices for access control rules, on page 2058

QoS rule components

QoS rules consist of several key components that control traffic rate limiting and filtering behavior. Understanding these components helps you configure effective quality of service policies.

QoS rules include these main components:

- State (Enabled/Disabled): Controls whether the rule is active and enforced by the system.
- Interfaces (Apply QoS On): Specifies which interfaces the rate limiting applies to.
- Traffic Limit Per Interface: Sets the rate limiting values for matching traffic.
- Conditions: Defines the traffic criteria that must be matched for the rule to apply.
- Comments: Provides documentation and change tracking for the rule.

State (Enabled/Disabled)

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Interfaces (apply QoS on)

You cannot save a QoS rule that rate limits all traffic. For each QoS rule, apply QoS on one of the options:

- Interfaces in Source Interface Objects—Rate limits traffic through the rule's source interfaces. If you choose this option, you must add at least one source interface constraint (cannot be **any**).
- Interfaces in Destination Interface Objects—Rate limits traffic through the rule's destination interfaces. If you choose this option, you must add at least one destination interface constraint (cannot be **any**).

Traffic limit per interface

A QoS rule enforces rate limiting *independently* on *each* interface specified with the Apply QoS On option. Aggregate rate limits cannot be applied to a group of interfaces.

You can rate limit traffic by Mbits per second. The default value of **Unlimited** prevents matching traffic from being rate limited.

You can independently rate limit download and upload traffic. The system determines download and upload directions based on the connection initiator.

If you specify a limit greater than the interface's maximum throughput, the system does not rate limit matching traffic. Maximum throughput may be affected by an interface's hardware configuration, which you specify in each device's properties (**Devices** > **Device Management**).

Conditions

Conditions specify the specific traffic that the rule handles. You can configure each rule with multiple conditions. Traffic must meet all conditions to match the rule. Each condition type has its own tab in the rule editor. For more information, see QoS rule conditions, on page 1700.

Comments

Each time you save changes to a rule, add comments. You might summarize the overall configuration for users, or note any changes made and the reasons.

In the policy editor, the system displays how many comments a rule has. In the rule editor, use the Comments tab to view existing comments and add new ones.

QoS rule conditions

Conditions specify the traffic the rule handles. You can configure each rule with multiple conditions. All conditions must be met for traffic to match the rule. Each condition type has its own tab in the rule editor. You can rate limit traffic using these conditions.

Related Topics

Interface rule conditions, on page 1700
Network rule conditions, on page 1701
User rule conditions, on page 1701
Application rule conditions, on page 1701
Best practice: use application-based rules over port-based rules, on page 1703
Recommendation: use URL conditions for website access control, on page 1704
Custom SGT rule conditions, on page 1704

Interface rule conditions

An interface rule condition is a network traffic control that

- controls traffic by its source and destination interfaces,
- uses predefined interface objects called security zones or interface groups to build interface conditions, and
- improves system performance by constraining rules to specific interfaces.

Interface rule condition characteristics

Interface objects segment your network by grouping interfaces across multiple devices to manage and classify traffic flow. For more information, refer to Interfaces, on page 2980.

QoS rules can be applied only on routed interfaces.



Note Constraining rules by interface is one of the best ways to improve system performance. If a rule excludes all of a device's interfaces, that rule does not affect that device's performance.

Network rule conditions

A network condition is a traffic matching criterion that

- controls or decrypts traffic based on source and destination IP addresses using inner packet headers,
- enables administrators to specify exact IP addresses or blocks for granular policy enforcement, and
- can be built using predefined network objects or by manually entering IP addresses or address blocks.

Network rule condition restrictions

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.



Note You *cannot* use FDQN network objects in identity rules.

User rule conditions

A user rule condition is a traffic matching mechanism that

- matches traffic based on whether the user initiating the connection or the group to which they belong
- can be configured for users in Microsoft Active Directory realms only, and
- enables policy enforcement based on user identity and group membership.

Special identities users

In addition to configuring users and groups for configured realms, you can set policies for the special identities users:

- Failed Authentication: User who failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.
- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified such as users who are not downloaded by a configured realm.

For access control rules only, you must first associate an identity policy with the access control policy as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

User rule condition example

Configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

Application rule conditions

Application rule conditions are application awareness and control mechanisms that

- identify and classify applications through IP traffic analysis
- enable control with organized filters
- require one enabled detector for each condition.

Application filters and control benefits

System-provided *application filters* help perform control by organizing applications according to basic characteristics, such as type, risk, business relevance, category, and tags. You can also create reuseable, user-defined filters using combinations of system-provided filters, or custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled, the system automatically enables all available system-provided detectors, or the most recently modified user-defined detector. For more information about application detectors, see Application detectors.

You can utilize both application filters and individually specified applications to ensure complete coverage. Understand the application's requirements to effectively order access control rules.

Application filters help you to quickly configure application control effectively. For example, you can easily use system-provided filters to create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

The system characterizes detected applications using several criteria which form effective filters.

Table 98: Application characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email applications are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Best practice: use application-based rules over port-based rules

Use application filtering criteria instead of port filtering to target traffic. Applications can be configured to use unique ports to bypass access control blocks.

Port conditions allow you to control traffic by its source and destination ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

The traditional method targets applications by using port specifications, but applications can manipulate unique ports to bypass these controls

Application filtering is recommended for applications to avoid blocking legitimate applications, like Firewall Threat Defense, that dynamically open separate channels for control and data flow. Port-based access control rules could block such applications, impacting desirable operations.

Using source and destination port constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP).

- For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).
- If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Port, protocol, and ICMP code rule conditions (concept)

A port, protocol, and ICMP code rule condition is a traffic matching criterion that

- matches traffic based on source and destination ports
- supports TCP, UDP, ICMP, and other protocols that do not use ports, and
- controls traffic flow in access control and other policy rules.

Port condition types

Port conditions match traffic based on the source and destination ports. Depending on the rule type, "port" can represent any of these types:

- TCP and UDP—Control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP (6) /22.
- ICMP—Control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—Control traffic using other protocols that do not use ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best practices for port-based rules

Specifying ports is the traditional method used to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Use application filtering criteria when possible; however, note that application filtering is unavailable in prefilter rules.

Use application filtering for FTP and similar applications that open separate channels for control and data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly and could result in blocking desirable connections.

Using source and destination port constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching non-TCP traffic with port conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although configuring port conditions can match non-TCP traffic, restrictions apply:

- Access control rules—For Classic devices, match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, add only network-based conditions: zone, IP address, port, and VLAN tag. The system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support only TCP port conditions.
- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies, sent in response to ICMP echo requests, are ignored. To match any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Recommendation: use URL conditions for website access control

Use URL conditions to control the websites that users on your network can access.

For complete information, see URL filtering with category and reputation, on page 2086.

Custom SGT rule conditions

A custom SGT rule condition is a traffic filtering mechanism that

- Use manually created SGT objects to filter traffic. Do not rely on ISE SGTs from a server.
- Correspond to the SGT attributes on the traffic to be controlled.
- Exclude user control considerations when using custom SGT traffic.

SGT traffic control without ISE configuration

If you do not configure ISE or ISE-PIC as an identity source, you can control traffic using Security Group Tags (SGTs) that were **not** assigned by ISE. SGTs specify traffic source privileges in a trusted network.

ISE SGT versus custom SGT rule conditions

Certain rules allow you to control traffic based on assigned SGT. The rule type and identity source configuration determine whether ISE-assigned SGTs or custom SGTs are used to match traffic with assigned SGT attributes.

Condition Type	Requires	SGTs Listed in Rule Editor
ISE SGT	ISE identity source	SGTs obtained by querying the ISE server, automatically updated metadata
Custom SGT	No ISE or ISE-PIC identity source	Static SGT objects you create



Note If using ISE SGTs to match traffic, any packet will match an ISE SGT rule if the SGT associated with the packet source IP address is recognized in ISE, even without an assigned SGT attribute.

Autotransition from custom Security Group Tags to ISE Security Group Tags

An autotransition from custom Security Group Tags (SGTs) to ISE Security Group Tags is a system process that automatically adjusts rule configuration when you configure ISE/ISE-PIC as an identity source after creating rules with custom Security Group Tags.

System behavior during autotransition

If you create rules that match custom SGTs, then configure ISE or ISE-PIC as an identity source, the system will:

- Disable **Security Group Tag** options: The system retains existing SGT objects, but you cannot modify them or add new ones.
- Retain existing rules: Existing rules with custom SGT conditions remain, but they do not match traffic, and you cannot add more custom SGT conditions.

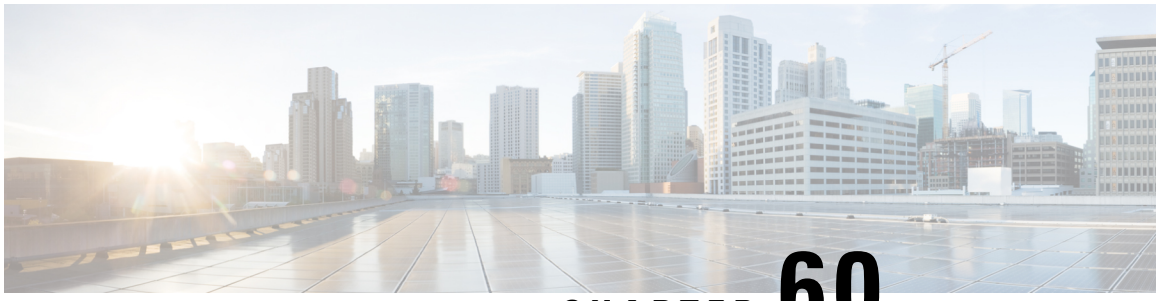
If you configure ISE, it is recommended to delete or disable existing rules with custom Security Group Tag conditions and use ISE attribute conditions to match traffic with Security Group Tag attributes.



PART XV

Secure Connections

- Secure Connections Overview, on page 1709
- Site-to-Site VPN, on page 1725
- SD-WAN, on page 1787
- Remote Access VPN, on page 1811
- Dynamic Access Policies , on page 1927
- VPN Monitoring and Troubleshooting in Security Cloud Control, on page 1943



CHAPTER 60

Secure Connections Overview

A virtual private network (VPN) connection establishes a secure tunnel between endpoints over a public network such as the Internet.

This chapter applies to Remote Access and Site-to-site VPNs on Secure Firewall Threat Defense devices. It describes the Internet Protocol Security (IPsec), the Internet Security Association and Key Management Protocol (ISAKMP, or IKE) and SSL standards that are used to build site-to-site and remote access VPNs.

- VPN types, on page 1709
- VPN basics, on page 1710
- VPN packet flow, on page 1713
- IPsec flow offload, on page 1713
- VPN licensing, on page 1714
- VPN encryption and performance, on page 1714
- Deprecated hash algorithms, encryption algorithms, and Diffie-Hellman modulus groups, on page 1720
- VPN topologies, on page 1720

VPN types

VPN types are network connection categories that

- provide secure, encrypted connections between remote locations and private networks
- support deployment models including remote access and site-to-site configurations, and
- use various protocols including SSL and IPsec for establishing secure tunnels.

Supported VPN connection types

The Cloud-Delivered Firewall Management Center supports these types of VPN connections:

- Remote Access VPNs in Firewall Threat Defense devices.

  Remote access VPNs provide secure, encrypted connections, or tunnels, between remote users and your company's private network. These connection use two devices: a VPN endpoint device, which is a workstation or mobile device with VPN client capabilities, and a VPN headend device, or secure gateway, at the edge of the corporate private network.

  Secure Firewall Threat Defense devices can be configured to support Remote Access VPNs over SSL or IPsec IKEv2 by the Cloud-Delivered Firewall Management Center. When acting as secure gateways,

these devices authenticate remote users, authorize access, and encrypt data to provide secure connections to your network. Only these devices support remote access VPN connections, managed by the Cloud-Delivered Firewall Management Center.

Secure Firewall Threat Defense secure gateways support the Secure Client full tunnel client. This client is required to provide secure SSL IPsec IKEv2 connections for remote users. This client automatically installs when a connection is established, so network administrators do not need to manually install or configure it on remote computers. It is the only client supported on endpoint devices.

- Site-to-site VPNs in Firewall Threat Defense devices.

  A site-to-site VPN connects networks in different geographic locations. You can create site-to-site IPsec connections between managed devices, and between managed devices and other Cisco or third-party peers. These peers can use either IPv4 and IPv6 addresses. Site-to-site tunnels are built using the Internet Protocol Security (IPsec) protocol suite and IKEv1 or IKEv2. After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel.

VPN basics

A VPN is a secure networking technology that

- uses tunneling to create secure connections between remote users and private corporate networks over public TCP/IP networks such as the Internet
- employs IPsec-based technologies with ISAKMP (IKE) and IPsec tunneling standards to build and manage tunnels, and
- enables bidirectional data transfer through tunnel endpoints that encapsulate and unencapsulate packets.

VPN tunnel management capabilities

ISAKMP and IPsec accomplish these tunnel management functions:

- Negotiate tunnel parameters.
- Establish tunnels.
- Authenticate users and data.
- Manage security keys.
- Encrypt and decrypt data.
- Manage data transfer across the tunnel.
- Manage data transfer inbound and outbound as a tunnel endpoint or router.

A device in a VPN functions as a bidirectional tunnel endpoint. It can receive plain packets from the private network, encapsulate them, create a tunnel, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets from the public network, unencapsulate them, and send them to their final destination on the private network.

After the site-to-site VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. A connection consists of the IP addresses

and hostnames of the two gateways, the subnets behind them, and the method the two gateways use to authenticate to each other.

VPN deployments use two primary device types:

- Hubs: Devices that enable secure VPN connectivity to and from one or more remote branch devices or spokes. Hubs also act as a gateway for spokes to communicate with each other.
- Spokes: Devices that connect over VPN to a hub to securely access the corporate resources behind the hub. Spokes communicate with each other through the hub.

Internet Key Exchange (IKE)

Internet Key Exchange (IKE) is a key management protocol that

- authenticates IPsec peers
- negotiates and distributes IPsec encryption keys, and
- automatically establishes IPsec security associations (SAs).

IKE negotiation phases and policies

The IKE negotiation comprises two phases:

- Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2.
- During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec.

Each phase uses proposals to negotiate a connection.

An IKE policy is a set of algorithms that two peers use to secure the IKE negotiation between them. IKE negotiation begins when each peer agrees on a common IKE policy. This policy defines security parameters to protect subsequent IKE negotiations. IKEv1 policies contain a single set of algorithms and a modulus group. In an IKEv2 policy, you can select multiple algorithms and modulus groups for peers to choose from during the Phase 1 negotiation. You may create a single IKE policy, or create multiple policies to give higher priority to preferred options. For site-to-site VPNs, you can create an IKE policy. IKEv1 and IKEv2 each support a maximum of 20 IKE policies, each with a different set of values. Assign a unique priority to each policy that you create. A lower priority number indocates a higher priority for the policy.

To define an IKE policy, specify:

- A unique priority (1 to 65,543, with 1 being the highest priority).
- An encryption method for the IKE negotiation to protect the data and ensure privacy.
- A Hashed Message Authentication Codes (HMAC) method (called integrity algorithm in IKEv2) to verify the sender's identity, and to confirm that the message is unchanged during transit.
- For IKEv2, a separate pseudorandom function (PRF) used as the algorithm to derive keying material and hashing operations required for the IKEv2 tunnel encryption. The options are the same as those used for the hash algorithm.
- A Diffie-Hellman group to determine the strength of the encryption key determination algorithm. The device uses this algorithm to derive the encryption and hash keys.

- An authentication method, to ensure the identity of the peers.
- A limit to the time the device uses an encryption key before replacing it.

When IKE negotiation begins, the initiating peer sends all its policies to the remote peer. The remote peer searches for a match with its own policies, in priority order. IKE policies match if both peers use the same encryption, hash (integrity and PRF for IKEv2), authentication, and Diffie-Hellman values. The SA lifetime must be less than or equal to the lifetime in the policy sent. If the lifetimes are not identical, the shorter value from the remote peer policy applies. By default, the Cloud-Delivered Firewall Management Center deploys an IKEv1 policy with the lowest priority for all VPN endpoints to ensure a successful negotiation.

IPsec

IPsec is a secure method for setting up a VPN that:

- provides data encryption at the IP packet level,
- transmits data over a public network through tunnels, which are secure, logical communication paths between two peers, and
- secures traffic that enters an IPsec tunnel by a combination of security protocols and algorithms.

With IPsec, data is transmitted over a public network through tunnels.

IPsec proposal policy components

An IPsec proposal policy defines the settings for IPsec tunnels. An IPsec proposal includes one or more crypto maps that are applied to the VPN interfaces in the devices. A crypto map combines all the components required to set up IPsec security associations, including:

- A proposal (or transform set) combines security protocols and algorithms that secure traffic in an IPsec tunnel. During the IPsec security association (SA) negotiation, peers search for a matching proposal. Apply the selected proposal to create an SA that protects data flows in the crypto map's access list and secures VPN traffic. There are separate IPsec proposals for IKEv1 and IKEv2. In IKEv1 proposals (or transform sets), you set one value for each parameter. In IKEv2 proposals, you can configure multiple encryption and integration algorithms for a single proposal.
- A crypto map combines all components required to set up IPsec security associations (SA), including IPsec rules, proposals, remote peers, and other parameters that are necessary to define an IPsec SA. When two peers try to establish an SA, they must each have at least one compatible crypto map entry.

  Dynamic crypto map policies are used in site-to-site VPNs when an unknown remote peer tries to start an IPsec security association with the local hub. The hub does not initiate the security association negotiation. Dynamic crypto policies allow remote peers to exchange IPsec traffic with a local hub, even if the hub does not know the remote peer's identity. A dynamic crypto map policy creates a crypto map entry before configuring all parameters. IPsec negotiation then dynamically configures the missing parameters to match a remote peer's requirements.

  Dynamic crypto map policies are applicable to both hub-and-spoke and point-to-point VPN topologies. To apply dynamic crypto map policies, specify a dynamic IP address for one of the peers in the topology and ensure that the dynamic crypto-map is enabled on this topology. In a full mesh VPN topology, you can apply only static crypto map policies.

Note Simultaneous IKEv2 dynamic crypto map is not supported for the same interface for both remote access and site-to-site VPNs in a Firewall Threat Defense device.

VPN packet flow

VPN packet flow is a security process that

- requires explicit permission through access control before allowing traffic to pass through,
- decrypts incoming tunnel packets before sending them to the Snort process,
- processes outgoing packets through Snort before encryption, and
- blocks tunnel traffic to the public source when the tunnel is down.

Access control requirements

Access control identifies the protected networks for each endpoint node of a VPN tunnel and determines which traffic is allowed to pass through the Firewall Threat Defense device and reach the endpoints. For remote access VPN traffic, a group policy filter or an access control rule must be configured to permit VPN traffic flow.

IPsec flow offload

IPsec flow offload is a performance optimization feature that

- offloads IPsec connections to field-programmable gate array (FPGA) or specialized hardware components after initial setup
- improves device performance by handling pre-decryption, decryption, pre-encryption, and encryption processing in hardware, and
- is enabled by default on supporting device models while allowing system software to handle inner flow security policies.

IPsec flow offload characteristics

After the initial setup of an IPsec site-to-site VPN or remote access VPN security association (SA), IPsec connections are offloaded to the field-programmable gate array (FPGA) in the device. This process improves device performance.In the Secure Firewall 1200 series, IPsec connections are offloaded to the Marvell Cryptographic Accelerator (CPT) to improve device performance.

Offloaded operations include pre-decryption and decryption processing on ingress and egress. The system software applies your security policies for traffic within the inner flow.

IPsec flow offload applies to these device types:

- Secure Firewall 1200

- Secure Firewall 3100
- Secure Firewall 4200

IPsec flow offload is also used when the device's VTI loopback interface is enabled.

For asymmetric flows in cluster distributed site-to-site VPN mode, IPsec flow offload lets the flow owner decrypt IPsec traffic in hardware that was forwarded over the cluster control link. This feature is not configurable and is always available with IPsec flow offload.

By default, the system enables IPsec flow offload on supported device models. To change the configuration, use FlexConfig to implement the **flow-offload-ipsec** command. For more information, refer to the ASA command reference.

VPN licensing

You do not need any specific license to enable VPN on a Firewall Threat Defense device, it is available by default.

The Cloud-Delivered Firewall Management Center determines whether to allow or block the usage of strong encryption in the Firewall Threat Defense device based on the export-controlled functionality in the device. You can enable this functionality when you register with the Cisco Smart License Manager. If you are using the evaluation license, or you did not enable export-controlled functionality, you cannot use strong encryption.

Choose **System** (⚙) > **Licenses** > **Smart Licenses** to verify this functionality in Cloud-Delivered Firewall Management Center.

If you created your VPN configurations with an evaluation license and later upgraded your license to a smart license with export-controlled functionality, check and update your encryption algorithms to use stronger encryption so that the VPNs function properly. Do not use DES-based encryption, as it is not supported.

VPN encryption and performance

When you configure VPN tunnel encryption, provide sufficient protection and maintain efficiency by balancing security and performance.

Because a VPN tunnel typically traverses a public network, most likely the Internet, you need to encrypt the connection to protect the traffic. You define the encryption and other security techniques with IKE policies and IPsec proposals. Using stronger tunnel encryption may lower system performance.

If your device license allows strong encryption, you can choose from a range of encryption and hash algorithms and Diffie-Hellman groups. This document does not provide specific guidance on which options you should choose. If you operate within a larger corporation or other organization, there might already be defined standards that you need to meet. If not, take the time to research the options.

Comply with security certification requirements

Review your certification requirements and the available options to plan your VPN configuration. Many VPN settings have options that allow you to comply with various security certification standards.

Decide encryption algorithms for VPN policies

When deciding which encryption algorithms to use for the IKE policy or IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN.

- For IKEv2, you can configure multiple encryption algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order.
- For IKEv1, you can select a single option only.
- For IPsec proposals, the algorithm is used by the Encapsulating Security Protocol (ESP), which provides authentication, encryption, and anti-replay services. ESP is IP protocol type 50. In IKEv1 IPsec proposals, the algorithm name is prefixed with ESP-.

If your device license qualifies for strong encryption, you can choose from the available encryption algorithms. If you are not qualified for strong encryption, you can select DES only.

Strong encryption license considerations

Note If you are qualified for strong encryption, before upgrading from the evaluation license to a smart license, check and update your encryption algorithms for stronger encryption so that the VPN configuration works properly. Choose AES-based algorithms. DES is not supported if you are registered using an account that supports strong encryption. After registration, you cannot deploy changes until you remove all uses of DES.

Available encryption algorithms

- AES-GCM—(IKEv2 only) Advanced Encryption Standard in Galois or Counter Mode is a block cipher mode that provides confidentiality and data-origin authentication. It offers greater security than AES.

  AES-GCM offers three different key strengths: 128-bit, 192-bit, and 256-bit keys. Longer keys increase security but reduce performance. NSA Suite B, a set of cryptographic algorithms that devices must support to meet federal cryptographic strength standards, requires GCM.
- AES—Advanced Encryption Standard is a symmetric cipher algorithm that provides greater security than DES and is computationally more efficient than 3DES. AES offers three different key strengths: 128-bit, 192-bit, and 256-bit keys. Longer keys increase security but reduce performance.
- DES—Data Encryption Standard, which encrypts using 56-bit keys, is a symmetric secret-key block algorithm. If your license account does not meet the requirements for export controls, this is your only option.
- Null, ESP-Null—A null encryption algorithm provides authentication without encryption. This method is not secure; use at your own discretion.

Decide which hash algorithms to use

In IKE policies, the hash algorithm creates a message digest to ensure your messages remain secure. In IKEv2, you can choose one hash algorithm for integrity and another for the pseudo-random function (PRF).

In IPsec proposals, the Encapsulating Security Protocol (ESP) uses the hash algorithm for authentication. In IKEv2 IPsec proposals, this is called the integrity hash. In IKEv1 IPsec proposals, the algorithm name has the ESP- prefix and an -HMAC suffix.

The system arranges your settings from the most secure to the least secure and negotiates with the peer in that order. For IKEv1, select only one option.

Select a hash algorithm that meets your security and performance needs:

- SHA (Secure Hash Algorithm)—The standard SHA (SHA1) produces a 160-bit digest.

  These SHA-2 options provide increased security and are available for IKEv2 configurations. Choose one if you require NSA Suite B cryptography compliance.

  - SHA256—Specifies the Secure Hash Algorithm SHA 2 with the 256-bit digest.
  - SHA384—Specifies the Secure Hash Algorithm SHA 2 with the 384-bit digest.
  - SHA512—Specifies the Secure Hash Algorithm SHA 2 with the 512-bit digest.

- Null or None (NULL, ESP-NONE)—(IPsec Proposals only) AUse a null hash algorithm only for testing purposes. If you select one of the AES-GCM options as the encryption algorithm, choose the null integrity algorithm. For these encryption standards, the integrity hash is ignored even if you choose a non-null option.

Decide which Diffie-Hellman modulus group to use

You can use Diffie-Hellman key derivation algorithms to generate IPsec security association (SA) keys. Each group has a different size modulus. A larger modulus increases security but requires more processing time. You must have a matching modulus group on both peers.

For AES encryption, use Diffie-Hellman (DH) Group 5 or higher to support the large key sizes required by AES. IKEv1 policies do not support all of the groups.

To implement the NSA Suite B cryptography specification, use IKEv2 and select one of the elliptic curve Diffie-Hellman (ECDH) options: 19, 20, or 21. Elliptic curve options and groups that use a 2048-bit modulus reduce exposure to attacks such as Logjam.

For IKEv2, you can configure multiple groups. The system orders the settings from the most secure to the least secure, then negotiates with the peer using that order. For IKEv1, you can select only a single option.

- 14—Diffie-Hellman Group 14: 2048-bit modular exponential (MODP) group. Considered good protection for 192-bit keys.
- 15—Diffie-Hellman Group 15: 3072-bit MODP group.
- 16—Diffie-Hellman Group 16: 4096-bit MODP group.
- 19—Diffie-Hellman Group 19: National Institute of Standards and Technology (NIST) 256-bit elliptic curve modulo a prime (ECP) group.
- 20—Diffie-Hellman Group 20: NIST 384-bit ECP group.
- 21—Diffie-Hellman Group 21: NIST 521-bit ECP group.
- 31—Diffie-Hellman Group 31: Curve25519 256-bit EC Group.

Decide VPN authentication methods

A VPN authentication method is a security mechanism that

- validates the identity of peers in a VPN connection
- enables secure communication between network devices, and
- ensures that only authorized devices can establish VPN connections.

Available authentication methods

VPNs support two primary authentication methods:

- **Preshared keys**: A secret key shared between two peers and used by IKE during the authentication phase. The same shared key must be configured at each peer or the IKE SA cannot be established.
- **Digital certificates**: Use RSA key pairs to sign and encrypt IKE key management messages. Certificates provide proof of communication between two peers.

VPN type support varies by authentication method.

Table 99: VPN authentication method support

VPN Type	Preshared Keys	Digital Certificates
Site-to-site IKEv1 and IKEv2	Supported	Supported
Remote Access (SSL and IPsec IKEv2)	Not supported	Supported

When you use digital certificate authentication, you need a Public Key Infrastructure (PKI) defined for peers to obtain digital certificates from a Certification Authority (CA). CAs manage certificate requests and issue certificates to participating network devices, providing centralized key management for all participating devices.

Preshared keys are difficult to manage for large networks. CAs make it easier to manage and scale your IPsec network. With a CA, you do not need to configure keys between all encrypting devices. Register each device with the CA to request its certificate. Devices with their own certificates and the CA's public key can authenticate other devices within the CA's domain.

Pre-shared keys

A pre-shared key is a secret key that

- enables you to share authentication credentials between two peers
- is used by IKE in the authentication phase, and
- must be configured identically on each peer or the IKE SA cannot be established.

Key configuration options

To configure pre-shared keys, select either a manual key or an automatically generated key. Specify this key in the IKEv1 or IKEv2 options. When you deploy your configuration, the key is configured on all devices in the topology.

PKI infrastructure and digital certificates

A PKI infrastructure is a centralized key management system that

- provides defined policies, procedures, and roles supporting public key cryptography
- generates, verifies, and revokes public key certificates, commonly known as digital certificates, and
- manages key pairs consisting of public and private keys for VPN endpoints to sign and encrypt messages.

Public key cryptography and certificate components

In public key cryptography, each endpoint of a connection has a key pair consisting of both a public and a private key. The key pairs are used by the VPN endpoints to sign and encrypt messages. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other, securing the data flowing over the connection.

Generate a general-purpose RSA, ECDSA, or EDDSA key pair for signing and encryption. Alternatively, generate separate key pairs for each purpose. Separate signing and encryption keys help to reduce key exposure. SSL uses a key for encryption, while IKE uses a key for signing. Separate keys for each purpose minimize exposure.

Certificates also ensure non-repudiation by providing proof that communication between two peers occurred.

You can obtain CA certificates by:

- Using the Simple Certificate Enrollment Protocol (SCEP) or Enrollment over Secure Transport (EST) to retrieve the CA's certificate from the CA server
- Manually copying the CA's certificate from another participating device

Trustpoints represent the object representation of a CA and associated certificates. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

A PKCS#12, or PFX, file holds the server certificate, any intermediate certificates, and the private key in one encrypted file. This type of file may be imported directly into a device to create a trustpoint.

A CA may also revoke certificates for peers that no longer participate in your network. Revoked certificates are either managed by an Online Certificate Status Protocol (OCSP) server or are listed in a certificate revocation list (CRL) stored on an LDAP server. A peer may check these before accepting a certificate from another peer.

Digital certificates or identity certificates

When you use digital certificates as the authentication method for VPN connections, peers are configured to obtain digital certificates from a Certificate Authority (CA). CAs are trusted authorities that sign certificates to verify their authenticity, thereby guaranteeing the identity of the device or user.

Digital certificates contain these components:

- The digital identification of the owner for authentication, such as name, serial number, company, department, or IP address.

- A public key needed to send and receive encrypted data to the certificate owner.
- The secure digital signature of a CA.

The certificate enrollment process includes these steps:

1. CA servers manage public CA certificate requests and issue certificates to participating network devices as part of a PKI.
2. Each participating device enrolls individually with a CA server, which validates identities and creates an identity certificate.
3. Each participating peer sends their identity certificate to the other peer to validate their identities and establish encrypted sessions with the public keys contained in the certificates.

Certificate enrollment

Using a PKI improves the manageability and scalability of your VPN because you do not have to configure preshared keys between all encrypting devices. Instead, enroll each participating device with a CA server, which validates identities and creates an identity certificate for the device. After completing enrollment, each participating peer sends its identity certificate to the other peer to validate their identities and establish encrypted sessions with the public keys contained in the certificates. For more information about enrolling Firewall Threat Defense device certificates, refer to Certificate enrollment objects, on page 2997.

Certificate authority certificates

In order to validate a peer’s certificate, each participating device must retrieve the CA's certificate from the server. A CA certificate is used to sign other certificates. It is self-signed and called a root certificate. This certificate contains the CA's public key, which you use to decrypt and validate the CA's digital signature and the contents of the received peer's certificate.

To obtain the CA certificate:

- Use the Simple Certificate Enrollment Protocol (SCEP) or Enrollment over Secure Transport (EST) to retrieve the CA’s certificate from the CA server
- Manually copy the CA's certificate from another participating device

Trustpoints

Once enrollment is complete, a trustpoint is created in the managed device. It is the object representation of a CA and associated certificates. A trustpoint includes the identity of the CA, CA-specific parameters, and one enrolled identity certificate.

PKCS#12 file

A PKCS#12 file, or PFX file holds the server certificate, intermediate certificates, and the private key in one encrypted file. This type of file may be imported directly into a device to create a trustpoint.

Revocation checking

A CA may also revoke certificates for peers that no longer participate in your network. Revoked certificates are either managed by an Online Certificate Status Protocol (OCSP) server or are listed in a certificate revocation list (CRL) stored on an LDAP server. A peer may check these before accepting a certificate from another peer.

Deprecated hash algorithms, encryption algorithms, and Diffie-Hellman modulus groups

Update your VPN configuration to use supported DH and encryption algorithms before upgrading to Firewall Threat Defense 6.70 or later.

- Update your IKE proposals and IPSec policies to match the ones supported in Firewall Threat Defense 6.70 or later.
- Deploy the configuration changes after updating to supported algorithms.

Support has been removed for these less secure ciphers from Firewall Threat Defense Version 6.70:

- **Diffie-Hellman GROUP 5** is deprecated for IKEv1 and IKEv2.
- Diffie-Hellman groups 2 and 24 have been removed.
- **Encryption algorithms**: 3DES, AES-GMAC, AES-GMAC-192, and AES-GMAC-256 have been removed.

Note **DES** continues to be supported in evaluation mode or for users who do not satisfy export controls for strong encryption.

NULL is removed in IKEv2 policy, but supported in both IKEv1 and IKEv2 IPsec transform-sets.

VPN topologies

When you create a new VPN topology you must give it a unique name, specify a topology type, and select the IKE version. You can select from three types of topologies, each containing a group of VPN tunnels:

- Point-to-point (PTP) topologies establish a VPN tunnel between two endpoints.
- Hub and spoke topologies establish a group of VPN tunnels connecting a hub endpoint to a group of spoke endpoints.
- Full mesh topologies establish a group of VPN tunnels among a set of endpoints.

There is no default pre-shared key for VPN authentication. You must define a pre-shared key manually or let the system generate it automatically. When choosing automatic, the Cloud-Delivered Firewall Management Center generates a pre-shared key and assigns it to all the nodes in the topology.

Point-to-point VPN topologies

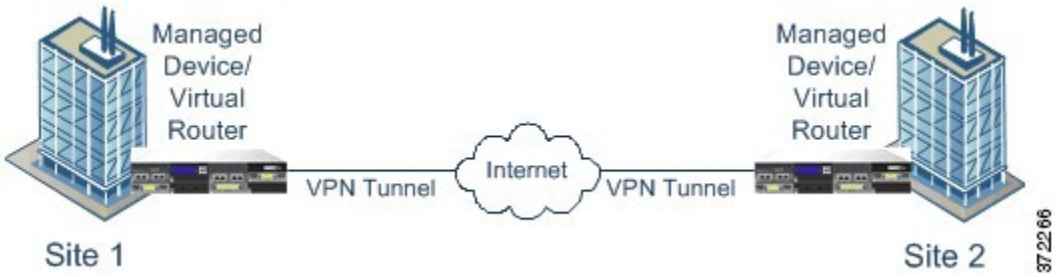
A point-to-point VPN topology is a network configuration that

- enables two endpoints to communicate directly with each other

- allows you to configure the two endpoints as peer devices, and
- permits either device to start the secured connection.

This diagram displays a typical point-to-point VPN topology.

Figure 365: Point-to-point VPN topology



Hub and spoke VPN topology

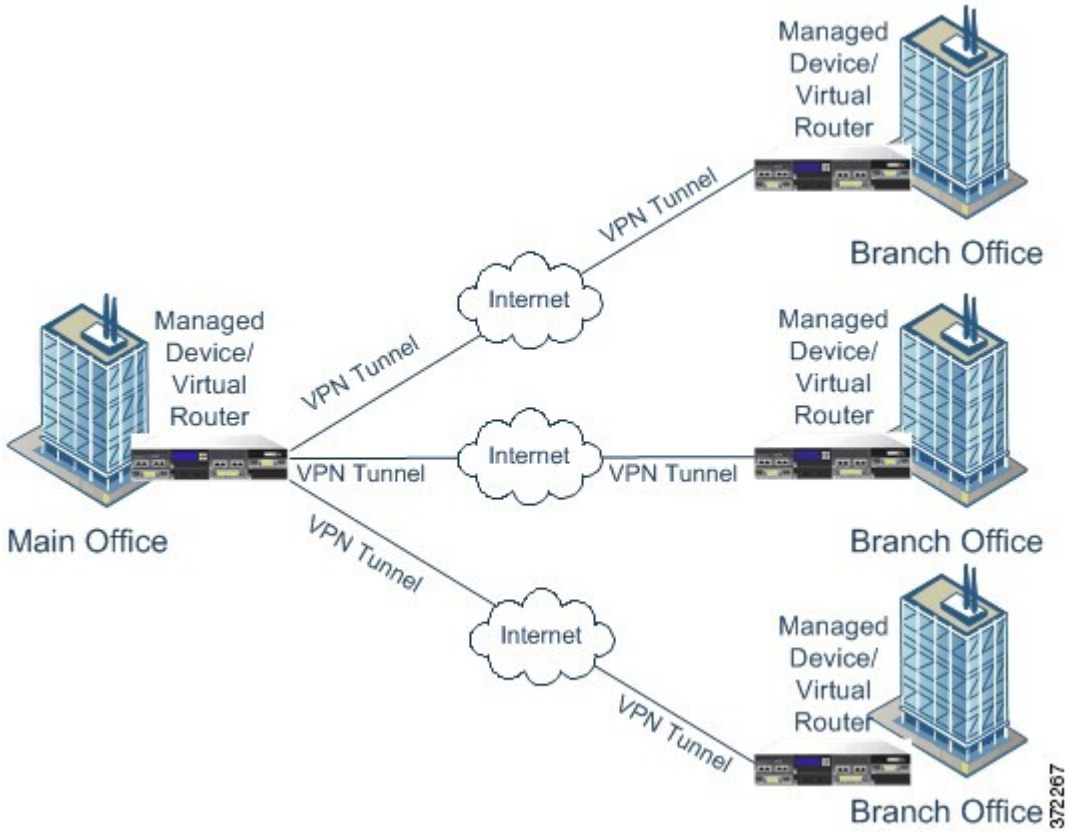
A hub and spoke VPN topology is a network architecture that

- connects a central endpoint (hub node) with multiple remote endpoints (spoke nodes)
- establishes each connection between the hub node and an individual spoke endpoint as a separate VPN tunnel, and
- enables hosts behind any of the spoke nodes to communicate with each other through the hub node.

The hub and spoke topology commonly represent a VPN that connects an organization's main and branch office locations using secure connections over the Internet or other third-party network. These deployments provide all employees with controlled access to the organization's network. Typically, the hub node is located at the main office. Spoke nodes are located at branch offices and start most of the traffic.

This diagram displays a typical hub and spoke VPN topology.

Figure 366: Hub and spoke VPN topology diagram



Full mesh VPN topology

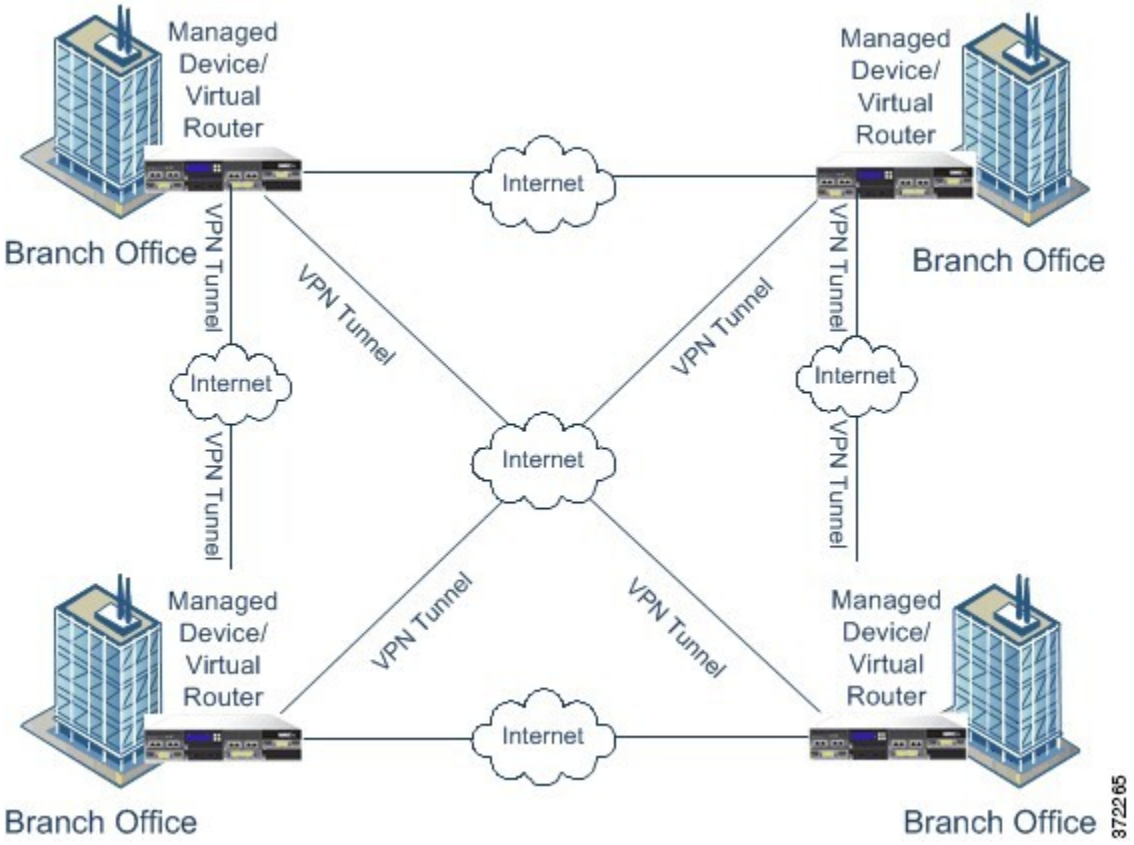
A full mesh VPN topology is a network configuration that

- allows all endpoints to communicate with every other endpoint by an individual VPN tunnel
- offers redundancy so that when one endpoint fails, the remaining endpoints can still communicate with each other, and
- commonly represents a VPN that connects a group of decentralized branch office locations.

The number of VPN-enabled managed devices you deploy in this configuration depends on the level of redundancy you require.

This diagram displays a typical full mesh VPN topology.

Figure 367: Full mesh VPN topology



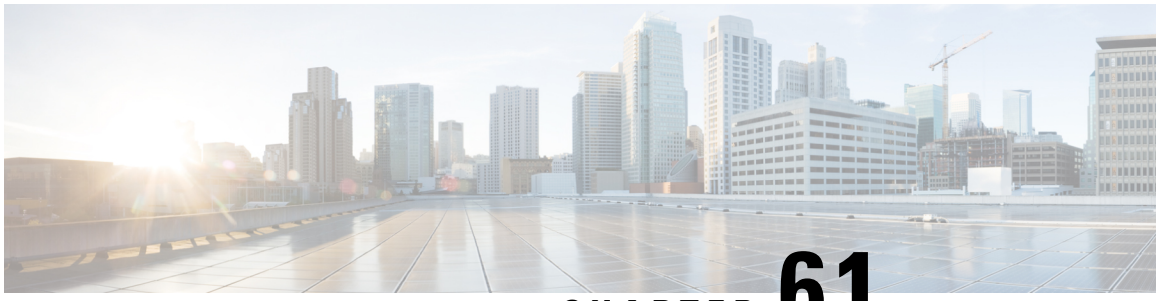
Implicit topologies

Implicit topologies are complex VPN network configurations that

- combine elements from the three main VPN topologies (full mesh, hub-and-spoke, and point-to-point)
- create more advanced network architectures than individual topology types, and
- provide customized connectivity solutions for specific network requirements.

Types of implicit topologies

- **Partial mesh**: A network where some devices use a full mesh topology, while others form either hub-and-spoke or point-to-point connections with some fully meshed devices. The partial mesh does not offer the redundancy found in a full mesh but costs less to implement. Peripheral networks use partial mesh topologies to connect to a fully meshed backbone.
- **Tiered hub-and-spoke**: A network of hub-and-spoke topologies where devices act as a hub in some topologies and as a spoke in others. Spoke groups can send traffic to their closest hub.
- **Joined hub-and-spoke**: A combination of two topologies (hub-and-spoke, point-to-point, or full mesh) that connect to form a point-to-point tunnel. For example, a joined hub-and-spoke topology could comprise two hub-and-spoke topologies, with the hubs acting as peer devices in a point-to-point topology.



CHAPTER 61

Site-to-Site VPN

- Site-to-site VPN, on page 1725
- Types of site-to-site VPN topologies, on page 1728
- License requirements for site-to-site VPN, on page 1728
- Prerequisites for site-to-site VPN, on page 1729
- Manage site-to-site VPNs, on page 1730
- Configure a policy-based site-to-site VPN, on page 1731
- Configure virtual tunnel interfaces, on page 1745
- Deploy a SASE tunnel on Umbrella, on page 1769
- Secure traffic using Cisco Secure Access and Firewall Threat Defense devices, on page 1776
- Monitor site-to-site topologies, on page 1781

Site-to-site VPN

A site-to-site VPN is a secure, encrypted VPN tunnel between two networks in different geographic locations. After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. You can create site-to-site IPsec connections between managed devices, and between managed devices and other Cisco or third-party peers.

Site-to-site VPN features

Secure Firewall Threat Defense site-to-site VPN supports these features:

- IPsec IKEv1 & IKEv2 protocols.
- Certificates and automatic or manual preshared keys for authentication.
- IPv4 & IPv6 addresses. All combinations of inside and outside are supported.
- Static and dynamic interfaces.
- HA environments for both Cloud-Delivered Firewall Management Center and Firewall Threat Defense.
- VPN alerts when the tunnel goes down and tunnel statistics.
- IKEv1 and IKEv2 back-up peer configuration for point-to-point extranet and hub-and-spoke VPNs.
- Extranet device as hub in 'Hub and Spokes' deployments.

- Dynamic IP address for a managed endpoint pairing with extranet device in 'Point to Point' deployments and for extranet device as an endpoint.

VPN topologies

VPN topology configuration requires specifying a unique name, a topology type, IKE version selection, and authentication method. The Cloud-Delivered Firewall Management Center configures site-to-site VPNs on Firewall Threat Defense devices only.

You can select from three types of topologies, containing one or more VPN tunnels:

- Point-to-point (PTP) deployments establish a VPN tunnel between two endpoints.
- Hub and Spoke deployments establish a group of VPN tunnels connecting a hub endpoint to a group of spoke nodes.
- Full Mesh deployments establish a group of VPN tunnels among a set of endpoints.

In the Cloud-Delivered Firewall Management Center, site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. Policies and proposals are sets of parameters that define the characteristics of a site-to-site VPN, such as the security protocols and algorithms that are used to secure traffic in an IPsec tunnel. Several policy types may be required to define a full configuration image that can be assigned to a VPN topology.

For authentication of VPN connections, configure a preshared key in the topology, or a trustpoint on each device. Preshared keys allow for a secret key, used during the IKE authentication phase, to be shared between two peers. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

IPsec and IKE

In the Cloud-Delivered Firewall Management Center, site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. Policies and proposals are sets of parameters that define the characteristics of a site-to-site VPN, such as the security protocols and algorithms that are used to secure traffic in an IPsec tunnel. Several policy types may be required to define a full configuration image that can be assigned to a VPN topology.

Authentication

For authentication of VPN connections, configure a preshared key in the topology, or a trustpoint on each device. Preshared keys allow for a secret key, used during the IKE authentication phase, to be shared between two peers. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

Extranet Devices

Each topology type can include extranet devices, devices that you don't manage in Cloud-Delivered Firewall Management Center. These include:

- Cisco devices that Cloud-Delivered Firewall Management Center supports, but for which your organization isn't responsible. Such as spokes in networks managed by other organizations within your company, or a connection to a service provider or partner's network.
- Non-Cisco devices. You can't use Cloud-Delivered Firewall Management Center to create and deploy configurations to non-Cisco devices.

Add non-Cisco devices, or Cisco devices not managed by the Cloud-Delivered Firewall Management Center, to a VPN topology as "Extranet" devices. Also specify the IP address of each remote device.

Guidelines for configuring site-to-site VPN in Secure Firewall Threat Defense devices

General guidelines

- Site-to-site VPN supports ECMP zone interfaces.
- Configure all nodes in a topology with either crypto ACL or a protected network. You cannot configure a topology with crypto ACL on one node and protected network on another.
- Configure a VPN connection across domains by using an extranet peer for the endpoint not in the current domain.
- You can backup Firewall Threat Defense VPNs using the Cloud-Delivered Firewall Management Center backup.
- Configure unique local IKE identity for all tunnels across all your VPN topologies.
- Ensure that IKE ports 500 and 4500 are not already in use and that no active PAT translations exist on those ports before configuring a site-to-site VPN. Configuring a site-to-site VPN on ports that are already in use will cause the service to fail to start.
- When configuring a site-to-site VPN between two devices managed by the same Cloud-Delivered Firewall Management Center, do not configure the devices as backup peers. Instead, configure one of the peer devices in the topology as an extranet device.
- If a spoke is part of multiple topologies with the same hub, for the Local Tunnel (IKE) Identity, we recommend that you use Key ID instead of an IP address.

  If the spoke uses an IP address as its local identity, the hub creates only one tunnel group. The spoke is then mapped to a specific virtual template interface, so the connection works with only one topology.

Limitations for configuring site-to-site VPN in Threat Defense devices

General limitations

- IKEv1 does not support CC/UCAPL-compliant devices. IKEv2 is recommended for these devices.
- VPN does not support network objects with a 'range' option.
- Firewall Threat Defense VPNs do not support PDF export and policy comparison.
- Tunnel status is not updated realtime, but at an interval of five minutes in the Cloud-Delivered Firewall Management Center.
- You cannot use the double quote character (") in pre-shared keys, replace it if already in use.

Crypto ACL limitations

- Cloud-Delivered Firewall Management Center supports only point-to-point VPN with crypto ACL and does not support tunnel health events.

- Cloud-Delivered Firewall Management Center does not verify the device interface address verification for transport mode when you select a crypto ACL.
- There is no support for automatic mirror ACE generation. Mirror ACE generation for the peer is a manual process on either side.

Topology management limitations

- You cannot move a VPN topology between domains.
- There is no per-tunnel or per-device edit option for Firewall Threat Defense VPNs, you can edit only the whole topology.

Types of site-to-site VPN topologies

This table provides an overview of the different site-to-site VPN topology options available for secure network connectivity configuration.

Site-to-Site VPN Topology	Description	More Information
SD-WAN Topology	Configure a secure branch network using the SD-WAN wizard. This wizard simplifies and automates VPN and routing configuration of your network using a hub and spoke topology, enabling SD-WAN capabilities.	SD-WAN wizard for secure branch network deployment, on page 1789
Route-Based VPN	Configure secure traffic dynamically between peers in a network based on routing over Virtual Tunnel Interfaces (VTIs).	Create a route-based site-to-site VPN, on page 1754
Policy-Based VPN	Configure secure traffic between peers in a network based on a static policy using protected networks.	Configure a policy-based site-to-site VPN, on page 1731
SASE Topology	Configure an IPsec IKEv2 tunnel from a Firewall Threat Defense device to Secure Access or an Umbrella Secure Internet Gateway (SIG). This tunnel forwards all internet-bound traffic to Secure Access or Umbrella SIG for inspection and filtering.	Secure traffic using Cisco Secure Access and Firewall Threat Defense devices, on page 1776 Configure a SASE tunnel for Umbrella, on page 1773

License requirements for site-to-site VPN

This topic lists the license requirements for setting up various types of site-to-site VPN connections.

License requirements for policy-based and route-based VPN

With the Essentials license, you can set up policy-based and route-based VPNs on your Firewall Threat Defense devices.

Depending on whether export-controlled functionality is enabled in your Smart License account, Cloud-Delivered Firewall Management Center determines whether to allow or block the usage of strong crypto on devices. To verify if export-controlled functionality is enabled for your Smart License account, choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Note If you use an evaluation license, or if you have not enabled the export-controlled functionality, you cannot use strong encryption for your VPN connections.

License requirements for deploying a SASE tunnel on Umbrella

To deploy SASE tunnels on Umbrella from Cloud-Delivered Firewall Management Center, you must enable your Smart License account with the export-controlled functionality. If this functionality is not enabled, you can only create a SASE topology, you cannot deploy tunnels on Umbrella.

Licenses for automatic tunnel between Secure Accessand Firewall Threat Defense devices

- Cloud-Delivered Firewall Management Center Essentials (formerly Base) license must allow export-controlled functionality. Choose **System** (⚙) > **Licenses** > **Smart Licenses** to verify this in Cloud-Delivered Firewall Management Center.

Note Evaluation license does not support this feature.

- Cisco Secure Access Essentials subscription license must be available.

Prerequisites for site-to-site VPN

Supported interfaces

Topology Type	Interface Type

Policy-Based	• Physical interfaces
	  • Non-management
	  • Interface Mode must be either Routed or None
	• Subinterface interfaces
	• Redundant interfaces
	• Etherchannel interfaces
	• VLAN interfaces

Route-Based	Static Virtual Tunnel Interfaces

Note The maximum number of supported security associations (SA) depends on the maximum VPN peers for your platform. For more information about maximum VPN peers, refer the data sheet of the Secure Firewall Threat Defense device.

General prerequisites

- You must have administrator privileges.
- You must be in a leaf domain to perform the site-to-site VPN configurations.

Manage site-to-site VPNs

You can view, create, modify, and deploy site-to-site VPN configurations and manage their deployment status.

The **Site-to-Site VPN** page provides a snapshot of site-to-site VPN tunnels. You can view the status of the tunnels and filter the tunnels based on the device, topology, or tunnel type. The page lists 20 topologies per page and you can navigate between pages to view more topology details. You can click individual VPN topologies to expand and view details of the endpoints.

Before you begin

For certificate authentication of your site-to-site VPNs, you must prepare the devices by allocating trustpoints as described in Certificates, on page 3079.

Procedure

Select **Devices** > **VPN** > **Site to Site** to manage your Firewall Threat Defense site-to-site VPN configurations and deployments.

The page lists the site-to-site VPNs topologies and indicates the status of tunnels using color codes:

- Active (Green)–There is an active IPsec Tunnel.
- Unknown (Amber)–No tunnel establishment event has been received from the device yet.
- Down (Red)–There are no active IPsec tunnels.
- Deployment Pending–Topology has not been deployed on the device yet.

Choose one of these actions:

- **Refresh**—View the updated status of the VPNs.
- **Add**—Create new policy based or route-based Site to Site VPNs.
- **Edit**—Modify the settings of an existing VPN topology.

  Note
  You cannot edit the topology type after you initially save it. To change the topology type, delete the topology and create a new one.

  Two users shouldn't edit the same topology simultaneously; however, the web interface doesn't prevent simultaneous editing.

- **Delete**—To delete a VPN deployment, click **Delete** ( ).
- Deploy—Choose **Deploy** > **Deploy**; see Deploy configuration changes, on page 202.

  Note
  Some VPN settings are validated only during deployment. Be sure to verify that your deployment was successful.

Configure a policy-based site-to-site VPN

Configure a policy-based site-to-site VPN to establish secure communication channels between network endpoints using crypto maps and specified network topologies.

Procedure

Step 1 Choose **Devices > Site To Site**. Then click **+ Site To Site VPN**, or edit a listed VPN topology.

Step 2 Enter a unique **Topology Name** and click the **Policy Based (Crypto Map)** radio button, and choose the **Network Topology** for this VPN.

Step 3 Choose the IKE versions to use during IKE negotiations. Check the **IKEv1** or **IKEv2** check box.

Default is IKEv2. Select either or both options as appropriate; select IKEv1 if any device in the topology doesn't support IKEv2.

You can also configure a backup peer for point-to-point extranet VPNs. For more information, see Firewall Threat Defense VPN endpoint options, on page 1732.

Step 4 Required: Add Endpoints for this VPN deployment by clicking **Add** (+) for each node in the topology.

Configure each endpoint field as described in Firewall Threat Defense VPN endpoint options, on page 1732.

- For Point to point, configure **Node A** and **Node B**.
- For Hub and Spoke, configure a **Hub Node** and **Spoke Nodes**
- For Full Mesh, configure multiple **Nodes**

Step 5 (Optional) Specify non-default IKE options for this deployment as described in Firewall Threat Defense VPN IKE options, on page 1736

Step 6 (Optional) Specify non-default IPsec options for this deployment as described in Firewall Threat Defense VPN IPsec options, on page 1739

Step 7 (Optional) Specify non-default Advanced options for this deployment as described in Firewall Threat Defense advanced site-to-site VPN deployment options, on page 1741 and click **Save**.

The policy-based site-to-site VPN is configured with the specified endpoints, IKE settings, and topology options.

What to do next

Deploy configuration changes.

Note Some VPN settings are validated only during deployment. Be sure to verify that your deployment was successful.

If you get an alert that your VPN tunnel is inactive even when the VPN session is up, follow the VPN troubleshooting instructions to verify and ensure that your VPN is active.

Firewall Threat Defense VPN endpoint options

Devices > Site To Site. Then click **+ Site To Site VPN**, or edit a listed VPN topology. Click the **Endpoints** tab.

Fields

Device

Choose an endpoint node for your deployment:

- A Firewall Threat Defense device managed by this Cloud-Delivered Firewall Management Center
- A Firewall Threat Defense high availability container managed by this Cloud-Delivered Firewall Management Center
- An **Extranet** device, any device (Cisco or third party) not managed by this Cloud-Delivered Firewall Management Center .

Device Name

For extranet devices only, provide a name for this device. We recommend naming it such that it is identifiable as an unmanaged device.

Interface

If you chose a managed device as your endpoint, choose an interface on that managed device.

For 'Point to Point' deployments, you can also configure an endpoint with dynamic interface. An endpoint with a dynamic interface can pair only with an extranet device and can't pair with an endpoint, which has a managed device.

You can configure device interfaces at **Devices** > **Device Management Add/Edit device** > **Interfaces**.

IP Address

- If you choose an extranet device, a device not managed by the Cloud-Delivered Firewall Management Center, specify an IP address for the endpoint.

  For an extranet device, select **Static** and specify an IP address or select **Dynamic** to allow dynamic extranet devices.

- If you chose a managed device as an endpoint, choose a single IPv4 address or multiple IPv6 addresses from the drop-down list. These IP addresses are already assigned to this interface on the managed device.

- All endpoints in a topology must have the same IP addressing scheme. IPv4 tunnels can carry IPv6 traffic and vice versa. The **Protected Networks** define which addressing scheme the tunneled traffic uses.

- If the managed device is a high-availability container, choose from a list of interfaces.

This IP is Private

Check the check box if the endpoint resides behind a firewall with network address translation (NAT).



Note Use this option only when the peer is managed by the same Cloud-Delivered Firewall Management Center and don't use this option if the peer is an extranet device.

Public IP address

If you checked the **This IP is Private** check box, specify a public IP address for the firewall. If the endpoint is a responder, specify this value.

Connection Type

Specify the allowed negotiation as bidirectional, answer-only, or originate-only. Supported combinations for the connection type are:

Table 100: Connection Type Supported Combinations

Remote Node	Central Node
Originate-Only	Answer-Only
Bi-Directional	Answer-Only

Remote Node	Central Node
Bi-Directional	Bi-Directional

Certificate Map

Choose a preconfigured certificate map object, or click **Add** (+) to add a certificate map object. The certificate map defines what information is necessary in the received client certificate to be valid for VPN connectivity. See Certificate map objects, on page 3058 for details.

Protected Networks



Caution Hub and Spoke topology—To avoid traffic drop for a dynamic crypto map, ensure that you don't select the protected network *any* for both the endpoints.

If the protected network is configured as *any*, on both the endpoints, the crypto ACL that works upon the tunnel is not generated.

Defines the networks that are protected by this VPN endpoint. Select the networks by selecting the list of Subnet/IP Address that define the networks that are protected by this endpoint. Click **Add** (+) to select from available Network Objects or add new Network Objects. See Create a network object, on page 2985 . Access control lists are generated from the choices made here.

- **Subnet/IP Address (Network)** —VPN endpoints can't have the same IP address and protected networks in a VPN endpoint pair cannot overlap. If protected networks for an endpoint contain IPv4 or IPv6 entries, the other endpoint's protected network must have at least one entry of the same type (IPv4 or IPv6). If it doesn't, the other endpoint's IP address must be of the same type and not overlap with the entries in the protected network. (Use /32 CIDR address blocks for IPv4 and /128 CIDR address blocks for IPv6.) If both of these checks fail, the endpoint pair is invalid.



Note By default, **Reverse Route Injection is enabled** check box is enabled in Cloud-Delivered Firewall Management Center.

Subnet/IP Address (Network) remains the default selection.

When you've selected Protected Networks as *Any* and observe default route traffic being dropped, disable the Reverse Route Injection. Choose **VPN** > **Site to Site** . Edit a VPN and then click **IPsec** > **Enable Reverse Route Injection** . Deploy the configuration changes to remove set reverse-route (Reverse Route Injection) from the crypto map configuration and remove the VPN-advertised reverse route that causes the reverse tunnel traffic to be dropped.

- **Access List (Extended)** —An extended access list provides the capability to control the type of traffic that will be accepted by this endpoint, like GRE or OSPF traffic. Traffic may be restricted either by address or port. Click **Add** (+) to add access control list objects.

Note Access Control List is supported only in the point to point topology.

Exempt VPN traffic from network address translation

Check this check box to exempt the VPN traffic from the Network Address Translation (NAT) rules.

If you do not exempt the VPN traffic from the NAT rules, the traffic gets dropped or is not routed through the VPN tunnel to the remote device. After you enable this option, you can view the NAT exemptions for the device in the NAT policy page (**Devices** > **NAT** , and click **NAT Exemptions**).

Inside interfaces directly connected to the internal network

Specify the security zone or interface group for the inside interface(s) where the protected networks reside. By default, the inside interface is any.

Click **+** to configure one or more interfaces from a security zone or an interface group that can map to one or more inside interfaces. Ensure that the interface type of the security zone or an interface group is Routed.

Advanced Settings

Enable Dynamic Reverse Route Injection —Reverse Route Injection (RRI) enables routes to be automatically inserted into the routing process, for the networks and hosts protected by a remote tunnel endpoint. Dynamic RRI routes are created only upon the successful establishment of IPsec security associations (SA's).

Note

- Dynamic RRI is supported only on IKEv2, and not supported on IKEv1 or IKEv1 + IKEv2.
- Dynamic RRI isn't supported on originate-only peer, Full Mesh topology, and Extranet peer.
- In Point-to-Point, only one peer can have dynamic RRI enabled.
- Between Hub and Spoke, only one of the endpoints can have dynamic RRI enabled.
- Dynamic RRI cannot be combined with a dynamic crypto map.

Send Local Identity to Peers —Select this option to send local identity information to the peer device. Select one of the following **Local Identity Configuration** from the list and configure the local identity:

- **IP address**—Use the IP address of the interface for the identity.
- **Auto**—Use the IP address for pre-shared key and Cert DN for certificate-based connections.
- **Email ID**—Specify the email ID to use for the identity. The email address can be up to 127 characters.
- **Hostname**—Use the fully qualified hostname.
- **Key ID**—Specify the key-id to use for the identity. The key ID must be fewer than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to the Cisco Umbrella Secure Internet Gateway (SIG).

For information about configuring a unique tunnel ID on Umbrella, see **Cisco Umbrella SIG User Guide**.

VPN Filter—Select an extended access list from the list or click **Add** to create a new extended access list object to filter the site-to-site VPN traffic.

The VPN filter provides more security and filters site-to-site VPN data using an extended access list. The extended access list object selected for the VPN filter lets you filter pre-encrypted traffic before entering the VPN tunnel and decrypted traffic that exits a VPN tunnel. The **sysopt permit-vpn** option, when enabled, would bypass the access control policy rules for the traffic coming from the VPN tunnel. When the **sysopt permit-vpn** option is enabled, the VPN filter helps in identifying and filtering the site-to-site VPN traffic.



Note The VPN filter is supported only on point to point, and hub and spoke topologies. It isn't supported on mesh topology.

For hub and spoke topology, you can choose to override the hub VPN filter on the spoke endpoints in case a different VPN filter needs to enabled on a specific tunnel.

Select the **Override VPN Filter on the Hub** option to override the hub VPN filter on the spokes. Select the **Remote VPN Filter** extended access list object or create an access list to override.



Note For an extranet device as a spoke, only the **Override VPN filter on the Hub** option is available.

For more information about sysopt permit-VPN, see Firewall Threat Defense advanced site-to-site VPN tunnel options , on page 1743 .

NAT Traversal (NAT-T)

NAT-T allow seamless communication between the peer Firewall Threat Defense devices when there are NAT devices between these devices. You cannot disable NAT Traversal (NAT-T) per VPN using the Cloud-Delivered Firewall Management Center UI. This option is available from Version 7.4.1. You can use FlexConfig to disable NAT-T. You can add the **crypto map** *map-name seq-num* **set nat-t-disable** command in the FlexConfig object. For example, **crypt map CSM_outside_map 1 set nat-t-disable**.



Note You must create the FlexConfig object with Deployment as 'EveryTime' and Type as 'Append'. This setting removes the command and adds it back for each deployment. If you configure Deployment as 'Once', then the command will be removed in the next deployment. You must be careful when you configure the FlexConfig object with this command. If Cloud-Delivered Firewall Management Center changes the sequence number, the deployment will fail.

Firewall Threat Defense VPN IKE options

For the versions of IKE you have chosen for this topology, specify the **IKEv1/IKEv2 Settings**.

Note Settings in this dialog apply to the entire topology, all tunnels, and all managed devices.

Devices > Site To Site. Then click **+ Site To Site VPN**, or edit a listed VPN topology. Click the **IKE** tab.

Fields

Policy

Choose the required IKEv1 or IKEv2 policy objects from the predefined list or create new objects to use. You can choose multiple IKEv1 and IKEv2 policies. IKEv1 and IKEv2 support a maximum of 20 IKE policies, each with a different set of values. Assign a unique priority to each policy that you create. The lower the priority number, the higher the priority.
We recommend that you do NOT use values such as 10, 20 and so on because the default IKEv2 policy for remote access VPN can have these values as its priority value. Before deployment, verify that the priority values of the IKE policies (site-to-site and remote access VPN) do NOT conflict.

Note When upgrading from Version 7.4.x or earlier—where the IKEv2 policy priority field is optional—to Version 7.6.x or later—where it is mandatory—leaving the priority field blank may lead to IKEv2 policy configuration issues and affect network traffic. To avoid these problems, ensure that a priority value is set for all IKEv2 policies either before or after the upgrade.

For details, see IKE policies, on page 3071

Authentication Type

Site-to-site VPN supports two authentication methods, pre-shared key and certificate. For an explanation of the two methods, see Decide VPN authentication methods, on page 1717.

Note In a VPN topology that supports IKEv1, the **Authentication Method** specified in the chosen IKEv1 Policy object becomes the default in the IKEv1 **Authentication Type** setting. These values must match, otherwise, your configuration will error.

- **Pre-shared Automatic Key**—The Cloud-Delivered Firewall Management Center automatically defines the pre-shared key for this VPN. Specify the **Pre-shared Key Length**, the number of characters in the key. The range is 1 to 27.

  The character " (double quote) isn't supported as part of pre-shared keys. If you've used " in a pre-shared key, ensure that you change the character after you upgrade to Secure Firewall Threat Defense 6.30 or higher.

- **Pre-shared Manual Key**—Manually assign the pre-shared key for this VPN. Specify the **Key** and then reenter the same to **Confirm Key**.

  When you choose this option for IKEv2, the **Enforce hex-based pre-shared key only** check box appears, check if desired. If enforced, you must enter a valid hex value for the key, an even number of 2-256 characters, using numerals 0-9, or A-F.

- **Certificate**—When you use certificates as the authentication method for VPN connections, peers obtain digital certificates from a CA server in your PKI infrastructure, and trade them to authenticate each other.

  In the **Certificate** field, select a preconfigured certificate enrollment object. This enrollment object generates a trustpoint with the same name on the managed device. The certificate enrollment object should be associated with and installed on the device, post which the enrollment process is complete, and then a trustpoint is created.

  A trustpoint is a representation of a CA or identity pair. A trustpoint includes the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate.

  Before you select this option, note these requirements:

  - Ensure you've enrolled a certificate enrollment object on all the endpoints in the topology—A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. Certificate Enrollment Objects are used to enroll your managed devices into your PKI infrastructure, and create trustpoints (CA objects) on devices that support VPN connections. For instructions on creating a certificate enrollment object, see Add certificate enrollment objects, on page 2999, and for instructions on enrolling the object on the endpoints see one of these as applicable:
    - Install a certificate using self-signed enrollment, on page 3084
    - Install a certificate using EST enrollment, on page 3084
    - Install a certificate using SCEP enrollment, on page 3085
    - Install a certificate using manual enrollment, on page 3086
    - Install a certificate using a PKCS12 file, on page 3087

Note For a site-to-site VPN topology, ensure that the same certificate enrollment object is enrolled in all the endpoints in the topology. For further details, see the table below.

- Refer to this table to understand the enrollment requirement for different scenarios. Some of the scenarios require you to override the certificate enrollment object for specific devices. See Manage object overrides, on page 2941 to understand how to override objects.

Certificate Enrollment Types	Device identity certificate for all endpoints is from the same CA		Device identity certificate for all endpoints is from different CAs
	Device-specific parameters are NOT specified in the certificate enrollment object	Device-specific parameters are specified in the certificate enrollment object	
Manual	No override required	Override required	Override required
EST	No override required	Override required	Override required

Certificate Enrollment Types	Device identity certificate for all endpoints is from the same CA		Device identity certificate for all endpoints is from different CAs
	Device-specific parameters are NOT specified in the certificate enrollment object	Device-specific parameters are specified in the certificate enrollment object	
SCEP	No override required	Override required	Override required
PKCS	Override required	Override required	Override required
Self-signed	NOT applicable	NOT applicable	NOT applicable

- Understand the VPN certificate limitations mentioned in Secure Firewall Threat Defense VPN certificate guidelines and limitations, on page 3079.

Note If you use a Windows Certificate Authority (CA), the default application policies extension is **IP security IKE intermediate**. If you use this default setting, you must select the **Ignore IPsec Key Usage** option in the Advanced Settings section on the **Key** tab in the **PKI Certificate Enrollment** dialog box for the object you select. Otherwise, the endpoints can't complete the site-to-site VPN connection.

Firewall Threat Defense VPN IPsec options

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure Endpoints for a Point to Point Topology and click the **IPsec** tab.

Note Settings in this dialog apply to the entire topology, all tunnels, and all managed devices.

Crypto-Map Type

A crypto map combines all the components required to set up IPsec security associations (SA). When two peers try to establish an SA, they must each have at least one compatible crypto map entry. IPsec security negotiation uses the proposals defined in the crypto map entry to protect the data flows specified by that crypto map's IPsec rules. Choose static or dynamic for this deployment's crypto-map:

- **Static**—Use a static crypto map in a point-to-point or full mesh VPN topology.
- **Dynamic**—Dynamic crypto-maps essentially create a crypto map entry without all the parameters configured. The missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a remote peer's requirements.

  Dynamic crypto map policies are applicable to both hub-and-spoke and point-to-point VPN topologies. To apply these policies, specify a dynamic IP address for one of the peers in the topology and ensure that the dynamic crypto-map is enabled on this topology. In a full mesh VPN topology, you can apply only static crypto map policies.

IKEv2 Mode

For IPsec IKEv2 only, specify the encapsulation mode for applying ESP encryption and authentication to the tunnel. This determines what part of the original IP packet has ESP applied.

- **Tunnel mode**—(default) Encapsulation mode is set to tunnel mode. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), hiding the ultimate source and destination addresses and becoming the payload in a new IP packet.

  The major advantage of tunnel mode is that you don't need to modify the end systems to receive the benefits of IPsec. This mode allows a network device, such as a router, to act as an IPsec proxy. That is, the router performs encryption on behalf of the hosts. The source router encrypts packets and forwards them along the IPsec tunnel. The destination router decrypts the original IP datagram and forwards it onto the destination system. Tunnel mode also protects against traffic analysis; with tunnel mode, an attacker can only determine the tunnel endpoints and not the true source and destination of the tunneled packets, even if they are the same as the tunnel endpoints.

- **Transport preferred**—Encapsulation mode is set to transport mode with an option to fall back to tunnel mode if the peer doesn't support it. In transport mode only the IP payload is encrypted, and the original IP headers are left intact. Therefore, the admin must select a protected network that matches the VPN interface IP address.

  This mode has the advantages of adding only a few bytes to each packet and allowing devices on the public network to see the final source and destination of the packet. With transport mode, you can enable special processing (for example, QoS) on the intermediate network based on the information in the IP header. However, the layer 4 header is encrypted, which limits examination of the packet.

- **Transport required**— Encapsulation mode is set to transport mode only, falling back to tunnel mode is allowed. If the endpoints can't successfully negotiate transport mode, due to one endpoint not supporting it, the VPN connection is not made.

Proposals

Click **Edit** ( ) to specify the proposals for your chosen IKEv1 or IKEv2 method. Select from the available **IKEv1 IPsec Proposals** or **IKEv2 IPsec Proposals** objects, or create and then select a new one. See Configure an IKEv1 IPsec proposal object, on page 3070 and Configure an IKEv2 IPsec proposal object, on page 3071 for details.

Enable Security Association (SA) Strength Enforcement

Enabling this option ensures that the encryption algorithm used by the child IPsec SA isn't stronger (in terms of the number of bits in the key) than the parent IKE SA.

Enable Reverse Route Injection

Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint.

Enable Perfect Forward Secrecy

Whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the Modulus Group list.

Modulus Group

The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing

time. The two peers must have a matching modulus group. For a full explanation of the options, see Decide which Diffie-Hellman modulus group to use, on page 1716.

Lifetime Duration
The number of seconds a security association exists before expiring. The default is 28,800 seconds.

Lifetime Size
The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires. The default is 4,608,000 kilobytes. Infinite data isn't allowed.

ESPv3 Settings

Validate incoming ICMP error messages
Choose whether to validate ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.

Enable 'Do Not Fragment' Policy
Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header.

Policy

- Copy DF bit—Maintains the DF bit.
- Clear DF bit—Ignores the DF bit.
- Set DF bit—Sets and uses the DF bit.

Enable Traffic Flow Confidentiality (TFC) Packets
Enable dummy TFC packets that mask the traffic profile which traverses the tunnel. Use the **Burst**, **Payload Size**, and **Timeout** parameters to generate random length packets at random intervals across the specified SA.

Note You can enable dummy Traffic Flow Confidentiality (TFC) packets at random lengths and intervals on an IPsec security association (SA). You must have an IKEv2 IPsec proposal set before enabling TFC.

Enabling TFC packets prevents the VPN tunnel from being idle. Thus the VPN idle timeout configured in the group policy doesn't work as expected when you enable the TFC packets.

Firewall Threat Defense advanced site-to-site VPN deployment options

The sections describe the advanced options you can specify in your site-to-site VPN deployment. These settings apply to the entire topology, all tunnels, and all managed devices.

Firewall Threat Defense VPN advanced IKE options

Advanced > IKE > ISAKAMP settings

IKE Keepalive

Enable or disables IKE Keepalives. You can set this option to EnableInfinite so that the device never starts the keepalive monitoring itself.

Threshold
Specifies the IKE keep alive confidence interval. This interval is the number of seconds allowing a peer to idle before beginning keepalive monitoring. The minimum and default interval is 10 seconds; the maximum interval is 3600 seconds.

Retry Interval
Specifies number of seconds to wait between IKE keep alive retries. The default is 2 seconds, the maximum is 10 seconds.

Identity Sent to Peers:
Choose the identity that the peers will use to identify themselves during IKE negotiations:

- **autoOrDN**(default)—Determines IKE negotiation by connection type: IP address for preshared key, or Cert DN for certificate authentication (not supported).
- **ipAddress**—Uses the IP addresses of the hosts exchanging ISAKMP identity information.
- **hostname**—Uses the fully qualified domain name of the hosts exchanging ISAKMP identity information. This name comprises the hostname and the domain name.



Note Enable or disable this option for all your VPN connections.

Peer Identity Validation
During IKE tunnel establishment, the peer provides its identity: either an IP address, a Fully Qualified Domain Name (FQDN), or a Distinguished Name (DN). It also presents a certificate, which contains none, some, or all of these fields.

If IKE peer identity validation is enabled, the Firewall Threat Defense compares the peer's identity to the respective field in the certificate to see if the information matches. If the information matches, then the peer's identity is validated and the Firewall Threat Defense establishes the tunnel. If the information does not match, the tunnel is not established.

- **Do not check**—Firewall Threat Defense does not validate the peer identity.
- **Required**—Firewall Threat Defense validates the peer identity.
- **If supported by cert**—Firewall Threat Defense validates the peer identity only if the peer provides a certificate.

Enable Aggressive Mode
Select this negotiation method for exchanging key information if the IP address isn't known and DNS resolution might not be available on the devices. Negotiation is based on hostname and domain name.

Enable Notification on Tunnel Disconnect
Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA does not match the traffic selectors for that SA. This notification is disabled by default.

Advanced > IKE > IKEv2 security association (SA) settings

More session controls are available for IKE v2 that limit the number of open SAs. By default, there's no limit to the number of open SAs.

Cookie Challenge
Whether to send cookie challenges to peer devices in response to SA initiate packets, which can help thwart denial of service (DoS) attacks. The default is to use cookie challenges when 50% of the available SAs are in negotiation. Select one of these options:

- Custom
- Never (default)
- Always

Threshold to Challenge Incoming Cookies
The percentage of the total allowed SAs that are in-negotiation. This triggers cookie challenges for any future SA negotiations. The range is zero to 100%.

Number of SAs Allowed in Negotiation
Limits the maximum number of SAs that can be in negotiation at any time. If used with Cookie Challenge, configure the cookie challenge threshold lower than this limit for an effective cross-check.

Maximum number of SAs Allowed
Limits the number of allowed IKEv2 connections. Default is unlimited.

Enable Notification on Tunnel Disconnect
Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA doesn't match the traffic selectors for that SA. Sending this notification is disabled by default.

Firewall Threat Defense VPN advanced IPsec options

Advanced > IPsec > IPsec Settings

The advanced IPsec settings include these configuration options:

- **Enable Fragmentation Before Encryption**: This option lets traffic travel across NAT devices that don't support IP fragmentation. It doesn't impede the operation of NAT devices that do support IP fragmentation.
- **Path Maximum Transmission Unit Aging**: Check to enable Path Maximum Transmission Unit (PMTU) Aging, the interval to reset the PMTU of a Security Association (SA).
- **Value Reset Interval**: Enter the number of minutes at which the PMTU value of an SA is reset to its original value. The range is 10 to 30 minutes, default is unlimited.

Firewall Threat Defense advanced site-to-site VPN tunnel options

This topic lists the advanced site-to-site VPN tunnel options.

Navigation path

Devices > Site To Site, then click **+ Site To Site VPN**, or edit a listed VPN Topology. Click the **Advanced** tab, and select **Tunnel** in the navigation pane.

Tunnel options

Only available for Hub and Spoke, and Full Mesh topologies. This section doesn’t appear for Point to Point configurations.

- **Enable Spoke to Spoke Connectivity through Hub**—Disabled by default. Choosing this field enables the devices on each end of the spokes to extend their connection through the hub node to the other device.

NAT settings

- **Keepalive Messages Traversal**—Enabled by default. This parameter is a global setting that enables NAT-T for all endpoints within a topology. Check this check box to enable keepalive messages for NAT traversal. NAT traversal keepalive is used for the transmission of keepalive messages when there's a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow.

  NAT traversal allows seamless communication between the peer Firewall Threat Defense devices when there are NAT devices between these devices. For hub and spoke topologies, this option is available only for the spoke.

  If you select this option, configure the **Interval**, in seconds, between the keepalive signals sent between the spoke and the middle device to indicate that the session is active. The value can be from 5 to 3600 seconds. The default is 20 seconds.

TrustSec (SGT) settings

- **Enable SGT propagation over Virtual Tunnel Interfaces**—Cisco TrustSec uses Security Group Tags (SGTs) to control access and enforce traffic on a network. This option enables SGT propagation over SVTIs and DVTIs of the VPN topology. To enable SGT propagation on a specific SVTI or DVTI, configure it in individual devices. Note that the Firewall Threat Defense device must be Version 10.0.0 and later.

Access control for VPN traffic

Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)—By default, the Firewall Threat Defense applies access control policy inspection on the decrypted traffic. Enable this option to bypass the ACL inspection. The Firewall Threat Defense still applies the VPN Filter ACL and authorization ACL downloaded from the AAA server to the VPN traffic.

Enable or disable the option for all your VPN connections. If you disable this option, ensure that the traffic is allowed by the access control policy or prefilter policy.



Note For route-based VPNs, **sysopt permit-vpn** does not work. You must always create access control rules to allow route-based VPN traffic.

Certificate map settings

- **Use the certificate map configured in the Endpoints to determine the tunnel**—If this option is enabled (checked), the tunnel is determined by matching the contents of the received certificate to the certificate map objects configured in the endpoint nodes.
- **Use the certificate OU field to determine the tunnel**—Indicates that if a node isn't determined based on the configured mapping (the above option) if selected, then use the value of the organizational unit (OU) in the subject distinguished name (DN) of the received certificate to determine the tunnel.

- **Use the IKE identity to determine the tunnel**—Indicates that if a node isn't determined based on a rule matching or taken from the OU (the above options) if selected, then the certificate-based IKE sessions are mapped to a tunnel based on the content of the phase1 IKE ID.
- **Use the peer IP address to determine the tunnel**—Indicates that if a tunnel isn't determined based on a rule matching or taken from the OU or IKE ID methods (the above options) if selected, then use the established peer IP address.

Configure virtual tunnel interfaces

Cloud-Delivered Firewall Management Center supports a routable logical interface called the Virtual Tunnel Interface (VTI).

Virtual tunnel interfaces

A virtual tunnel interface is a routable logical interface that

- does not require static mapping of IPsec sessions to a physical interface
- associates the IPsec tunnel endpoint with a virtual interface, and
- supports static and dynamic routing policies.

Virtual tunnel interface capabilities

Cloud-Delivered Firewall Management Center supports VTIs as an alternative to policy-based VPN. VTIs support route-based VPN with IPsec profiles attached to the end of each tunnel. VTIs use static or dynamic routes. The device encrypts or decrypts the traffic from or to the tunnel interface and forwards it according to the routing table. Deployments become easier, and having VTI which supports route-based VPN with dynamic routing protocol also satisfies many requirements of a virtual private cloud. Cloud-Delivered Firewall Management Center enables you to easily migrate from crypto-map based VPN configuration to VTI-based VPN.

You can configure route-based VPN with static or dynamic VTI using the site-to-site VPN wizard. Traffic is encrypted using static route, BGP, OSPFv2/v3, or EIGRP.

You can create a routed security zone, add VTI interfaces to it, and define access control rules for the decrypted traffic control over the VTI tunnel.

There are two types of VTI interfaces: static VTI and dynamic VTI.

You can create VTI-based VPNs between:

- Two Firewall Threat Defense devices.
- A Firewall Threat Defense and public cloud.
- One Firewall Threat Defense and another Firewall Threat Defense with service provider redundancy.
- A Firewall Threat Defense and any other device with VTI interfaces.
- A Firewall Threat Defense and another device with policy-based VPN configuration.

For more information, see Static VTI, on page 1746 and Dynamic VTI, on page 1747.

Static VTI

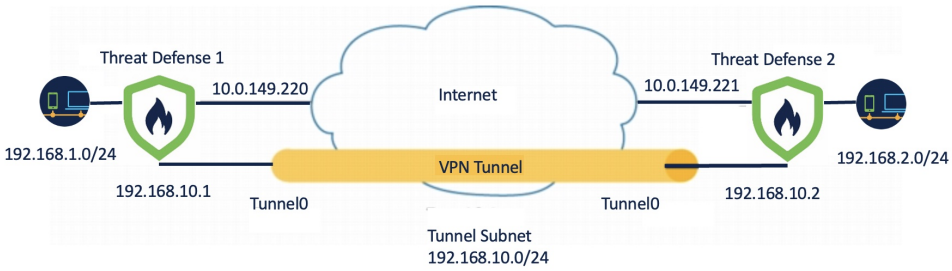
A static VTI is a virtual tunnel interface that

- uses tunnel interfaces to create a tunnel that is always-on between two sites
- requires defining a physical interface as a tunnel source, and
- supports a maximum of 1024 VTIs per device.

To create a static VTI interface in the Management Center, see Add a VTI interface, on page 1752.

The figure shows a VPN topology using static VTIs.

Figure 368: Static VTI VPN topology



On Threat Defense 1:

- Static VTI IP address is 192.168.10.1
- Tunnel source is 10.0.149.220
- Tunnel destination is 10.0.149.221

On Threat Defense 2:

- Static VTI IP address is 192.168.10.2
- Tunnel source is 10.0.149.221
- Tunnel destination is 10.0.149.220

Benefits of SVTI

Static VTI provides these benefits:

- Minimizes and simplifies configuration.

  You do not have to track all remote subnets for a crypto map access list, and configure complex access lists or crypto maps.

- Provides a routable interface.

  Supports IP routing protocols such as BGP, EIGRP, and OSPFv2/v3, and static routes.

- Supports backup VPN tunnels
- Supports load balancing using ECMP.

- Supports virtual routers.
- Provides differential access control for VPN traffic.

  You can configure a VTI with a security zone and use it in an AC policy. This configuration:

  - Allows you to classify and differentiate VPN traffic from clear-text traffic and permit VPN traffic selectively.
  - Provides differential access-control for VPN traffic across different VPN tunnels.

Dynamic VTI

A dynamic VTI is a VPN interface that

- uses a virtual template for dynamic instantiation and management of IPsec interfaces
- dynamically generates a unique virtual access interface for each VPN session
- supports multiple IPsec security associations and accepts multiple IPsec selectors proposed by the spoke.

Benefits of dynamic VTI

- Minimizes and simplifies configuration.

  You do not have to configure complex access lists or crypto maps.
- Simplifies management.
  - Easily manage peer configuration for large enterprise hub and spoke deployments.
  - Use only one dynamic VTI for multiple spokes, instead of configuring one static VTI per spoke.
- Provides a routable interface.

  Supports IP routing protocols such as BGP, EIGRP, and OSPFv2/v3.
- Simplifies scaling

  Addition of new spokes does not require any additional VPN configuration on the hub. You may need to update NAT and routing configurations depending upon the setup.
- Support for backup VPN tunnels.
- Supports dynamic spokes.

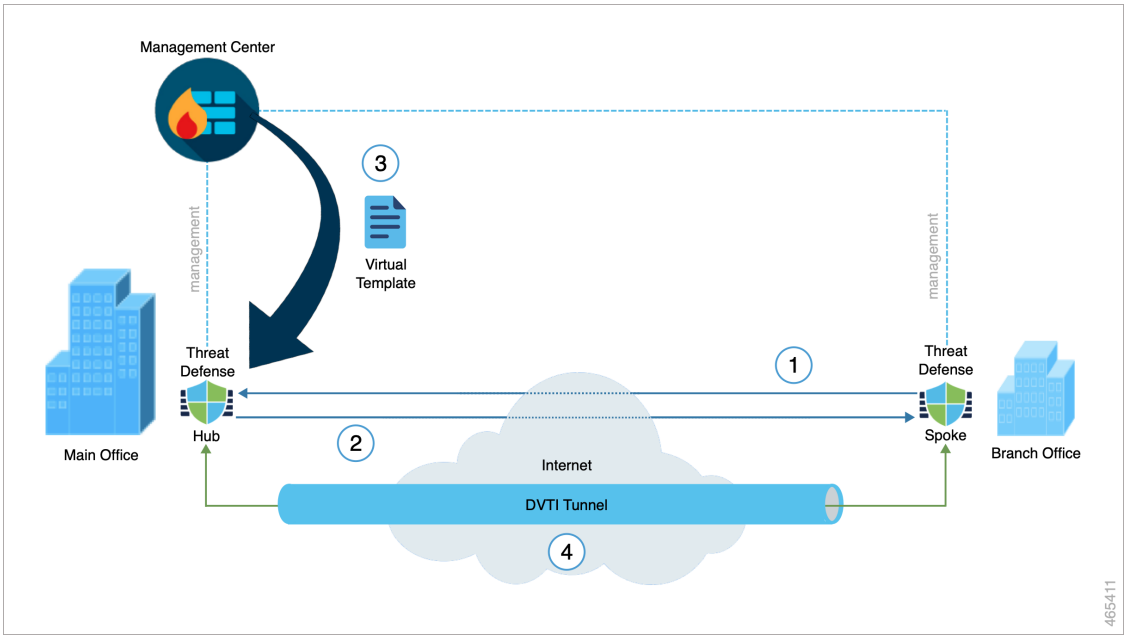
  You do not have to update the hub configuration for spoke's DHCP IP address changes.
- Conserves IP addresses.
  - Uses the IP unnumbered interface functionality to borrow the IP address from another physical or loopback interface.
  - All virtual access interfaces associated to a dynamic VTI use the same IP address.
- Supports virtual routers.
- Provides differential access control for VPN traffic.

You can configure a VTI with a security zone and use it in an AC policy. This configuration:

- Allows you to classify and differentiate VPN traffic from clear-text traffic and permit VPN traffic selectively.
- Provides differential access-control for VPN traffic across different VPN tunnels.

How Does the Cloud-Delivered Firewall Management Center Create a Dynamic VTI Tunnel for a VPN Session

When a spoke initiates a tunnel request with the hub:



1. The spoke initiates an IKE exchange with the hub for a VPN connection.
2. The hub authenticates the spoke.
3. The Cloud-Delivered Firewall Management Center assigns a dynamic virtual template on the hub for the spoke.

   The virtual template dynamically generates a virtual access interface on the hub. This interface is unique for the VPN session with the spoke.
4. The hub establishes a dynamic VTI tunnel with the spoke using the virtual access interface.
   a. The hub and spoke exchange traffic over the tunnel using:
      - Specific traffic proposed by the spokes over IKE exchanges.
      - BGP/OSPF/EIRGP protocols over the IPsec tunnel.
   b. After the VPN session ends, the tunnel disconnects and the hub deletes the corresponding virtual access interface.

To create a dynamic VTI interface in the Cloud-Delivered Firewall Management Center, see Add a VTI interface, on page 1752.

To configure a route-based site-to-site VPN using dynamic VTI, see Configure dynamic VTI for a route-based site-to-site VPN, on page 1765.

Virtual Routers and Dynamic VTI

You can create virtual routers, associate dynamic VTIs with these virtual routers, and extend the capabilities of dynamic VTIs in your network. You can associate dynamic VTIs either with global or user-defined virtual routers. You can assign a dynamic VTI to only one virtual router.

A virtual router associated with:

- A dynamic VTI is called an Indoor VRF (IVRF).
- A tunnel source interface is known as Front Door VRF (FVRF).

A dynamic VTI and its corresponding protected network interface must be part of the same virtual router. You must map the borrow IP interface and the dynamic VTI to the same virtual router. A tunnel source interface can be part of multiple virtual routers.

To configure virtual routers using dynamic VTI for a route-based site-to-site VPN, refer to Configure a virtual router with dynamic VTI, on page 1765.

For more information about a configuration example, refer to Secure traffic from networks with multiple virtual routers over a site-to-site VPN with dynamic VTI, on page 1409

Guidelines for virtual tunnel interfaces

General configuration guidelines

- VTIs are only configurable in IPsec mode.
- Dynamic VTI supports only Firewall Threat Defense devices from Version 7.3.
- You can use static, BGP, EIGRP IPv4, OSPFv2/v3 routes for traffic using the tunnel interface.
- In an HA configuration with dynamic routing, the standby device cannot access the known subnets through the VTI tunnels as these tunnels are created with the active IP address.
- For static and dynamic VTI, ensure that you don't use the borrow IP interface as the tunnel source IP address for any VTI interface.
- When you configure a route-based site-to-site VPN using static or dynamic VTI interfaces, ensure that the value of the TTL hop is more than one if you use BGP.
- If NAT has to be applied, the IKE and ESP packets are encapsulated in the UDP header.
- Ensure that you remove unused transform sets from route-based IKE VPN configurations. If you have unused transform sets, the VPN tunnels may not come up and encapsulation and decapsulation will not happen. Use **show run crypto** and **show VPN-sessiondb detail l2l** commands to view the transform sets.

IPv6 address guidelines

- VTI supports IPv6 addresses.
- You can use an IPv6 address for the tunnel source interface and use the same address as the tunnel endpoint.

- The Cloud-Delivered Firewall Management Center supports the following combinations of VTI IP (or internal networks IP version) over public IP versions: IPv6 over IPv6, IPv4 over IPv6, IPv4 over IPv4, and IPv6 over IPv4.
- VTI supports static and dynamic IPv6 addresses as the tunnel source and destination.
- The tunnel source interface can have IPv6 addresses and you can specify the tunnel endpoint address. If you don't specify the address, by default, the device uses the first IPv6 global address in the list as the tunnel endpoint.

IKE support and guidelines

- Static VTI supports IKE v1 and v2.
- Dynamic VTI supports only IKE version v2, and uses IPsec for sending and receiving data between the tunnel's source and destination.
- IKE and IPsec security associations are re-keyed continuously regardless of data traffic in the tunnel. This ensures that VTI tunnels are always up.
- Tunnel group name must match what the peer sends as its IKEv1 or IKEv2 identity.
- For IKEv1 in LAN-to-LAN tunnel groups, you can use names which aren't IP addresses, if the tunnel authentication method is digital certificates and/or the peer is configured to use aggressive mode.

MTU behavior

- MTU for VTIs is automatically set based on the underlying physical interface.
- For dynamic VTI, the virtual access interface inherits MTU from the configured tunnel source interface. If you don't specify the tunnel source interface, the virtual access interface inherits the MTU from the source interface from which the device accepts the VPN session request.

VTI scaling and guidelines

- Configure a maximum of 1024 static and dynamic VTIs in a device.

  While calculating the VTI count, consider nameif subinterfaces, portchannel configurations, and platform VLAN limitations.
- If you configure more than 400 VTIs in a device in a high-availability setup, configure 45 seconds as the unit holdtime.

Routing protocol support

VTI supports multiple routing protocols for different IP versions.

- VTI supports IPv6 BGP.
- VTI supports IPv4 EIGRP.
- VTI supports IPv4 and IPv6 OSPF.

ECMP configuration guidelines

- Configure spokes' static VTIs in an ECMP zone to load balance application traffic. If you do not configure the ECMP zone, the remaining paths act as backup paths when the primary path goes down.

- You can associate VTI interfaces with ECMP zones and configure ECMP static routes to achieve these benefits:
  - Load balancing (Active/Active VTIs)—Connection can flow over any of the parallel VTI tunnels.
  - Seamless connection migration—When a VTI tunnel becomes unreachable, the flows are seamlessly migrated to another VTI interface that is configured in the same zone.
  - Asymmetric routing—Forward traffic flow through one VTI interface and configure the reverse traffic flow through another VTI interface.

  For information on configuring ECMP, see ECMP routing methods, on page 1431.

Access control rule guidelines

- By default, all traffic sent through a VTI is encrypted.
- Access control rules can be applied on a VTI interface to control traffic through VTI.
- For route-based VPNs, Bypass Access Control policy for decrypted traffic (**sysopt connection permit-VPN**) does not work. You must always create access control rules to allow route-based VPN traffic.

Crypto map guidelines

- If you use dynamic crypto maps and dynamic VTIs in your site-to-site VPNs, only the dynamic VTI tunnels will come up. This behaviour occurs because both the crypto maps and dynamic VTIs try to use the default tunnel group. We recommend that you migrate your site-to-site VPNs to dynamic VTIs or use static crypto maps with their own tunnel-groups.
- VTI and crypto map configurations can coexist on the same physical interface if the peer address configured in the crypto map and the tunnel destination for the VTI are different.

Topology guidelines

- Dynamic VTI supports only the hub and spoke topology in the management center.
- We recommend that you configure only a single hub for a route-based hub and spoke topology. To configure a topology with multiple hubs for a set of spokes, with one hub as the backup hub, configure multiple topologies with a single hub and the same set of spokes.

Dynamic VTI and virtual router guidelines

- A dynamic VTI and its corresponding protected network interface must be part of the same virtual router.
- You must map the borrow IP interface and the dynamic VTI to the same virtual router.
- User-defined virtual routers support only BGPv4/v6 and OSPFv2 routing protocols.
- A tunnel source interface can be in a different user-defined virtual router than that associated with the dynamic VTI.

Firewall mode guideline

Configure VTI only in routed mode.

Limitations for Virtual Tunnel Interfaces

This topic lists the limitations for configuring Virtual Tunnel Interfaces in Management Center.

General limitations

- VTIs do not support clustering.
- VTI tunnel interfaces are not supported for SNMP or Syslog in Platform Settings.

Static VTI limitations

- Only 20 unique IPSec profiles are supported.
- In policy-based routing, you can configure VTI only as an egress interface.
- You cannot configure a VTI interface as a network interface for a remote access VPN policy.

Dynamic VTI limitations

- Dynamic VTI does not support ECMP , VRF in multi-instance, Clustering, IKEv1, or QoS.
- If a spoke has a dynamic IP address and a hub has a dynamic VTI behind a NAT, the tunnel status will be unknown.
- For a dynamic extranet, when multiple spokes establish a connection, the site-to-site monitoring dashboard does not show the individual tunnels.
- If you configure a hub with dynamic VTI behind NAT with dynamic spokes, the VPN monitoring data will not be accurate.
- You cannot configure a VTI interface as a network interface for a remote access VPN policy.

Backup VTI limitations

- Flow resiliency across tunnel failovers isn't supported. For example, the clear text TCP connection gets lost after a tunnel failover, and you need to reinitiate any FTP transfer that took place during the failover.
- Certificate authentication isn't supported in backup VTI.

Add a VTI interface

For configuring a route-based site-to-site VPN, you must create a VTI interface on the devices at both the nodes of the VTI tunnel.

When you specify the tunnel type as dynamic and configure the related parameters, the Cloud-Delivered Firewall Management Center generates a dynamic virtual template. The virtual template dynamically generates the virtual access interface that is unique for each VPN session.

Before you begin

Configure a loopback interface for redundancy of static and dynamic VTI VPN tunnels. For more information, see Configure a loopback interface, on page 1053.

For a Secure Firewall 1200/3100/4200device, IPsec flow offload is also used when the device's VTI loopback interface is enabled.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the **Edit** icon next to the device on which you want to create a VTI interface.

Step 3 Choose **Add Interfaces > Virtual Tunnel Interface**.

In the **Add Virtual Tunnel Interface** dialog box, configure these parameters:

a) Select the **Tunnel Type** as **Static** or **Dynamic**, and enter the name and description for the interface. By default, the interface is enabled.

b) (Optional) Choose a security zone from the **Security Zone** drop-down list to add the static VTI or dynamic VTI interface to that zone. In the **Priority** field, enter the priority to load balance the traffic across multiple VTIs.

If you want to perform traffic inspection based on a security zone, add the VTI interface to the security zone and configure an access control (AC) rule. To permit the VPN traffic over the tunnel, you need to add an AC rule with this security zone as the source zone.

The priority range is from 0 to 65535. The lowest number has the highest priority. This option is not applicable for dynamic VTI.

c)

d) Depending on the tunnel type, do one of the following:

- For a dynamic VTI, enter a unique ID in the range of 1 to 10413 in the **Template ID** field.
- For a static VTI, enter a unique tunnel ID in the range of 0 to 10413 in the **Tunnel ID** field.

e) (Optional for dynamic VTI) Choose the tunnel source interface from the **Tunnel Source** drop-down list.

The VPN tunnel terminates at this interface, a physical or loopback interface. Choose the IP address of the interface from the drop-down list. You can select the IP address irrespective of the IPsec tunnel mode. In case of multiple IPv6 addresses, select the address that you want to use as the tunnel endpoint.

f) Under **IPSec Tunnel Mode**, click the **IPv4** or **IPv6** radio button to specify the traffic type over the IPsec tunnel.

g) Under **IP Address**:

- **Configure IP**: Enter the IPv4 or IPv6 address for the static VTI interface. You cannot configure an IP address for a dynamic VTI interface. Use the **Borrow IP** field for the dynamic VTI interface.
- **Borrow IP (IP unnumbered)**: Choose a physical or loopback interface from the drop-down list, the VTI interface inherits this IP address.

  Ensure that you use an IP address different from the tunnel source IP address. You can use this option for a static or dynamic VTI interface.

  Click Add + to configure a loopback interface. The loopback interface helps to overcome path failures. If an interface goes down, you can access all interfaces through the IP address assigned to the loopback interface.

Step 4 Click **OK**.

Step 5 Click **Save**.

Create a route-based site-to-site VPN

You can configure a route-based site-to-site VPN for the following two topologies:

- **Point to Point**: Configure VTIs on both nodes of the tunnel and use the wizard to configure the VPN.
- **Hub and Spoke**: Configure VTIs on the hub and the spokes. Configure the hub with a dynamic VTI and spokes with static VTIs.

You can configure an extranet device as the hub and managed devices as spokes. You can configure multiple hubs and spokes, and also configure backup hubs and spokes.

- For extranet hubs and spokes, you can configure multiple IPs as backup.
- For managed spokes, you can configure a backup static VTI interface along with the primary VTI interface.

For more information on VTI, see Virtual tunnel interfaces, on page 1745.

Note All references to VTI stands for static VTI and dynamic VTI, unless mentioned.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.

Step 2 Click **+ Site To Site VPN**.

Step 3 Enter a name for the VPN topology in the **Topology Name** field.

Step 4 Choose **Route Based (VTI)** and do one of the following:

- Select **Point to Point** as the network topology. To configure endpoints for a route-based **Point to Point** topology, refer to Configure endpoints for a point-to-point VPN topology, on page 1755.
- Select **Hub and Spoke** as the network topology. To configure endpoints for a route-based **Hub and Spoke** topology, refer to Configure endpoints for a hub and spoke topology, on page 1757.

Step 5 (Optional) Specify the **IKE**, **IPsec**, and **Advanced** options and click **Save**.

For more information, refer to these topics:

- Firewall Threat Defense VPN IKE options, on page 1736
- Firewall Threat Defense VPN IPsec options, on page 1739
- Firewall Threat Defense advanced site-to-site VPN deployment options, on page 1741

What to do next

After you configure VTI interfaces and VTI tunnel on both the devices, you must configure:

- A routing policy to route the VTI traffic between the devices over the VTI tunnel. For more information, see Configure routing and AC policies for VTI, on page 1766.
- An access control rule to allow encrypted traffic. Choose **Policies** > **Access Control**.

Configure endpoints for a point-to-point VPN topology

Configure these parameters to configure endpoints for a route-based site-to-site VPN for the **Point to Point** topology nodes:

Before you begin

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Create a route-based site-to-site VPN, on page 1754 and click the **Endpoints** tab.

Procedure

Step 1 Under **Node A**, in the **Device** drop-down menu, select the name of the registered device (Firewall Threat Defense) or extranet as the first endpoint of your VTI tunnel.

For an extranet peer, specify the following parameters:

a. Specify the name of the device.

b. Enter the primary IP address in the **Endpoint IP address** field. If you configure a backup VTI, add a comma and, specify the backup IP address.

c. Click **OK** .

After configuring the above parameters for the extranet hub, specify the pre-shared key for the extranet in the **IKE** tab.

Note
The AWS VPC has **AES-GCM-NULL-SHA-LATEST** as the default policy. If the remote peer connects to AWS VPC, select **AES-GCM-NULL-SHA-LATEST** from the **Policy** drop-down list to establish the VPN connection without changing the default value in AWS.

Step 2 For a registered device, you can specify the VTI interface for Node A from the **Virtual Tunnel Interface** drop-down list.

The selected tunnel interface is the source interface for Node A and the tunnel destination for Node B.

If you want to create a new interface on Node A, click the Add + icon and configure the fields as described in Add a VTI interface, on page 1752 .

If you want to edit the configuration of an existing VTI, select the VTI in the **Virtual Tunnel Interface** drop-down field and click **Edit VTI**.

Step 3 If your Node A device is behind a NAT device, check the **Tunnel Source IP is Private** check box. In the **Tunnel Source Public IP Address** field, enter the tunnel source public IP address.

Step 4 **Send Local Identity to Peers** —Select this option to send local identity information to the peer device. Select one of the following **Local Identity Configuration** from the list and configure the local identity:

- **IP address** —Use the IP address of the interface for the identity.
- **Auto** —Use the IP address for pre-shared key and Cert DN for certificate-based connections.
- **Email ID** —Specify the email ID to use for the identity. The email ID can be up to 127 characters.
- **Hostname** —Use the fully qualified hostname.
- **Key ID** —Specify the key-id to use for the identity. The key ID must be fewer than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to a Cisco Umbrella Secure Internet Gateway (SIG).

For information about configuring a unique tunnel ID on Umbrella, see **Cisco Umbrella SIG User Guide** .

Step 5 (Optional) Click **Add Backup VTI** to specify an extra VTI as the backup interface and configure the parameters.

Note
Ensure that both peers of the topology do not have the same tunnel source for the backup VTI. A device cannot have two VTIs with the same tunnel source and tunnel destination; hence, configure a unique tunnel source and tunnel destination combination.

Though the virtual tunnel interface is specified under Backup VTI, the routing configuration determines which tunnel to be used as primary or backup.

Step 6 Expand **Advance Settings** to configure additional configurations for the device. For more information, see Configure advanced settings for a point-to-point topology in a route-based VPN, on page 1757 .

Step 7 Repeat the above procedure for Node B and click **OK** .

What to do next

- (Optional) Specify the **IKE** options for the deployment as described in Firewall Threat Defense VPN IKE options, on page 1736.
- (Optional) Specify the **IPsec** options for the deployment as described in Firewall Threat Defense VPN IPsec options, on page 1739.
- (Optional) Specify the **Advanced** options for the deployment as described in Firewall Threat Defense advanced site-to-site VPN deployment options, on page 1741.
- Click **Save**.
- To route traffic to the VTI, choose **Devices** > **Device Management** , edit the threat defense device and click the **Routing** tab.

  You can configure the static routes, or use BGP, OSPF v2/v3, or EIGRP for routing the VPN traffic.
- To permit VPN traffic, choose **Policies** > **Access Control heading** > **Access Control**. Add a rule specifying the security zone of the VTI. For a backup VTI, ensure that you include the backup VTI in the same security zone as that of the primary VTI.

Configure advanced settings for a point-to-point topology in a route-based VPN

Configure advanced settings to enhance the functionality and security of a point-to-point topology in a route-based VPN beyond the basic configuration parameters.

Before you begin

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure endpoints for a point-to-point VPN topology, on page 1755 and expand **Advance Settings**.

Procedure

Step 1 Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.

Step 2 Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the spokes and peers.

Step 3 Choose one of the following from the **Connection Type** drop-down list:

Answer Only: The device can only respond when a peer device initiates a connection, it can't initiate any connection.

Bidirectional: The device can initiate or respond to a connection. This is the default option.

Configure endpoints for a hub and spoke topology

You can create a route-based site-to-site VPN using dynamic VTI only for hub and spoke topologies. The hub can use only a dynamic VTI and the spokes can use only static VTI interfaces. You can also configure an extranet device as a hub.

Configure these parameters to configure endpoints for a route-based site-to-site VPN for the **Hub and Spoke** topology nodes:

Before you begin

Configure the basic parameters for a hub and spoke topology in a route-based VPN as described in Create a route-based site-to-site VPN, on page 1754 and click the **Endpoints** tab.

Procedure

Step 1 Under **Hub Nodes**:

a) Click Add + to configure the hub node in the **Add Endpoint** dialog box.

b) Choose a hub from the **Device** drop-down list.

Note
A Firewall Threat Defense device running on software version 7.2 can't be configured as a hub. It must be an extranet or a device running on software version 7.3 or later.

For an extranet hub, configure these parameters:

1. Enter the name of the device.

2. Enter the primary IP address. If you configure a backup VTI, add a comma, and then specify the backup IP address.
3. Click **OK**.

After configuring the above parameters for the extranet hub, specify the pre-shared key for the extranet in the **IKE** tab.

Note
The AWS VPC has **AES-GCM-NULL-SHA-LATEST** as the default policy. If the remote peer connects to AWS VPC, select **AES-GCM-NULL-SHA-LATEST** from the **Policy** drop-down list to establish the VPN connection without changing the default value in AWS.

c) Choose a dynamic VTI from the **Dynamic Virtual Tunnel Interface** drop-down list.

Tunnel source configuration is mandatory for a dynamic VTI as the management center needs this information to determine the tunnel destination of the spoke.

Click Add + to add a new dynamic VTI. We recommend that you configure the Borrow IP for the dynamic interface from a loopback interface.

If you want to edit an existing dynamic VTI, select the interface, and click **Edit VTI**.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box and configure the tunnel source IP address in the **Tunnel Source Public IP Address** field.
e) Click **Routing Policy** to configure the routing policy for the hub.
f) Click **AC Policy** to configure the access control policy.
g) Expand **Advance Settings** to configure additional configurations on the hub. For more information, see Configure advanced settings for hub and spokes in a route-based VPN, on page 1760 .
h) Click **OK**.

Step 2 Under **Spoke Nodes**:

a) Click Add + to configure the spoke in the **Add Endpoint** dialog box.
b) Choose a spoke from the **Device** drop-down list.

For an extranet spoke, configure these parameters:

1. Enter the name of the device.
2. In **Endpoint IP Address**, choose one of these options:
   - **Static**: Enter the IP address of the device, and the backup IP address, if required.
   - **Dynamic**: Choose this option to dynamically assign the IP addresses for the extranet spokes.
3. Click **OK**.

c) Choose a static VTI from the **Static Virtual Tunnel Interface** drop-down list.

Click Add + to add a new static VTI. Tunnel IP of the static VTI is auto populated, ensure that this IP address is unique for the spoke.

If you want to edit an existing static VTI, select the interface, and click **Edit VTI**.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box. The management center needs the tunnel source interface address to configure the tunnel destination

IP address on the spokes. In the **Tunnel Source Public IP Address** field, enter the tunnel source public IP address.

e) (Optional) **Send Local Identity to Peers** —Check this check box to send the local identity information to the peer device. Choose one of the following parameters from the **Local Identity Configuration** drop-down list and configure the local identity:

- **IP address**—Use the IP address of the interface for the identity.
- **Auto**—Use the IP address for pre-shared key and Cert DN for certificate-based connections.
- **Email ID**—Specify the email ID to use for the identity. The email ID can be up to 127 characters.
- **Hostname**—Use the fully qualified hostname.
- **Key ID**—Specify the key-id to use for the identity. The key ID must be less than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to the Cisco Umbrella Secure internet Gateway (SIG).

For more information about configuring a unique tunnel ID on Umbrella, see *Cisco Umbrella SIG User Guide* .

f) (Optional) Click **Add Backup VTI** to specify an extra VTI interface as the backup interface.

Note
Ensure that both peers of the topology do not have backup VTI configured on the same tunnel source. For instance, if Peer A is having two VTIs (primary and a backup) configured with a single tunnel source interface, say, 10.10.10.1/30, then Peer B also can't have its 2 VTIs with a single tunnel source IP, say 20.20.20.1/30.

Though the virtual tunnel interface is specified under backup VTI, the routing configuration determines which tunnel to be used as primary or backup.

g) Click **Routing Policy** to configure the routing policy for the spoke.
h) Click **AC Policy** to configure the access control policy.
i) Expand **Advance Settings** to configure additional configurations on the spoke and click **OK** . For more information, see Configure advanced settings for hub and spokes in a route-based VPN, on page 1760.

What to do next

1. (Optional) Specify the **IKE** options for the deployment as described in Firewall Threat Defense VPN IKE options, on page 1736.
2. (Optional) Specify the **IPsec** options for the deployment as described in Firewall Threat Defense VPN IPsec options, on page 1739.
3. (Optional) Specify the **Advanced** options for the deployment as described in Firewall Threat Defense advanced site-to-site VPN deployment options, on page 1741.
4. Click **Save**.

Configure advanced settings for hub and spokes in a route-based VPN

After configuring basic parameters for a hub and spoke in a route-based VPN, you can configure additional advanced settings to optimize the VPN topology's functionality and performance.

Before you begin

Configure the basic parameters for a hub and spoke in a route-based VPN as described in Configure endpoints for a hub and spoke topology, on page 1757 and expand **Advance Settings**.



Note Only the **Connection Type** field is applicable to the device with Version 7.2. The remaining fields don't apply to this version of the device.

Procedure

Step 1 Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.

For a hub, you must check this check box if you use BGP as the routing protocol. This configuration ensures that the loopback IP address is shared in the BGP routing table.

For a spoke, this option is enabled by default.

Step 2 Add the **Protected Networks** to define the networks protected by the VPN endpoint. Click Add + to select a protected network.

For a hub, configure the protected networks behind the hub. This information and the spoke's protected network generate the spoke access list.

You cannot create a static route for a virtual access interface on a hub with dynamic VTI. The hub creates and deletes these interfaces dynamically during tunnel establishment and termination.

For a spoke, configure the spoke's protected network.

To enable static routing for the spokes, after you configure the endpoints for your topology, click the **IPsec** tab and check the **Enable Reverse Route Injection** check box.

You do not need this option if you use BGP, OSPF, or EIGRP.

Step 3 Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the spokes and peers.

For a hub: During an IKE exchange, the hub advertises the dynamically created virtual access interfaces to the spokes, and the spokes advertise their VTI IP addresses to the hub.

For a spoke: By default, this option is enabled.

Step 4 In the **Connection Type** drop-down list, choose one of the following:

Answer Only: The device can only respond when a peer device initiates a connection, it can't initiate any connection.

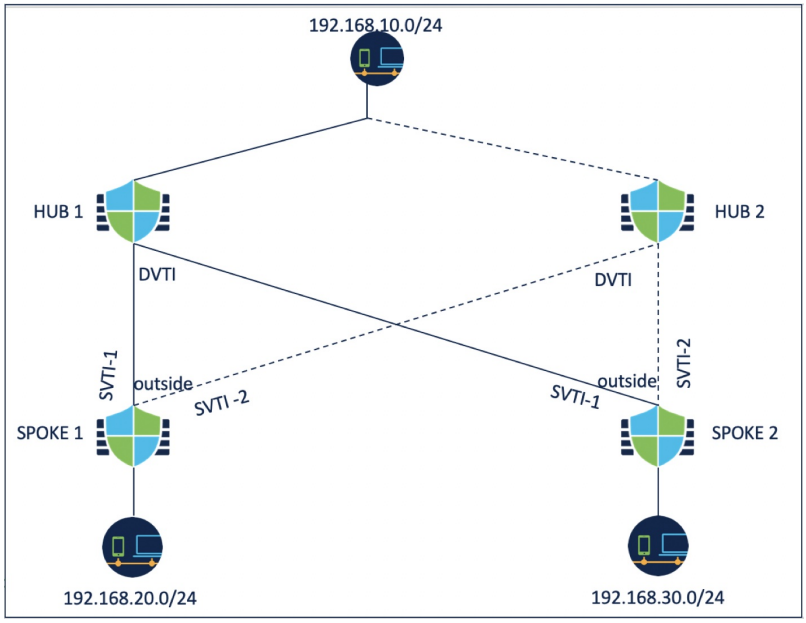
Bidirectional: The device can initiate or respond to a connection. This is the default option.

Configure multiple hubs in a route-based VPN

You can configure a topology with multiple hubs for a set of spokes. With one hub as the backup hub, you can configure multiple topologies with a single hub and the same set of spokes.

In this example, there are two hubs connected to the same set of spokes. Hub 1 is the primary hub and Hub 2 is the secondary hub. To configure this network in the Cloud-Delivered Firewall Management Center, you must configure two route-based hub and spoke topologies:

- Topology 1: Hub 1 connected to spoke 1 and spoke 2.
- Topology 2: Hub 2 connected to spoke 1 and spoke 2.



To configure topology 1:

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site** and click **+ Site To Site VPN**.

Step 2 Enter a name for the VPN topology in the **Topology Name** field.

Step 3 Choose **Route Based (VTI)** and do one of these topologies:

- Select **Point to Point** as the network topology. To configure endpoints for a route-based **Point to Point** topology, see Configure endpoints for a point-to-point VPN topology, on page 1755.
- Select **Hub and Spoke** as the network topology. To configure endpoints for a route-based **Hub and Spoke** topology, see Configure endpoints for a hub and spoke topology, on page 1757.

Step 4 Configure the IKE version and click the **Endpoints** tab.

Step 5 Under **Hub Nodes**:

a) Click Add + to add the hub.
b) Choose Hub 1 from the **Device** drop-down list.
c) Choose a dynamic VTI from the **Dynamic Virtual Tunnel Interface** drop-down list or click Add + to add a new dynamic VTI.

We recommend that you configure the Borrow IP for the dynamic interface from a loopback interface.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box and configure the tunnel source IP address in the **Tunnel Source Public IP Address** field.
e) Click **Routing Policy** to configure the routing policy for the hub. You can configure dynamic routing using BGP.
f) Expand **Advance Settings**. You can configure these advanced settings for the hub to enable IKEv2 routing, which can be used if you do not use dynamic routing.

- (Optional) Check the **Send Virtual Tunnel Interface IP to the peers** check box.
- Check the **Allow incoming IKEv2 routes from the peers** check box for the hub to accept routes from the spokes and update the routing table.
- Choose **Connection Type** as Bidirectional from the drop-down list.

g) Click **OK**.

Step 6 Under **Spoke Nodes**:

a) Click Add + to add a spoke.
b) Choose Spoke 1 from the **Device** drop-down list.
c) Choose SVTI-1 as the static VTI for the spoke from the **Static Virtual Tunnel Interface** drop-down list or click Add + to add a new static VTI.

Choose the outside interface as the tunnel source of SVTI-1. Tunnel IP of SVTI-1 is autopopulated, ensure that this IP address is unique for spoke 1 across peers in both the topologies.

d) Expand **Advance Settings**. If you do not use dynamic routing, you can configure these settings to enable IKEv2 routing for the spoke.

- Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.
- Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the peers.
- Choose **Connection Type** as Bidirectional from the drop-down list.

e) Click **OK**.
f) Repeat steps 5a to 5e to add spoke 2. Configure SVTI-1 as the static VTI of spoke 2.

Step 7 Configure the IKE and IPSec parameters as required or use the default values.

What to do next

1. Configure topology 2 with hub 2, spoke 1, and spoke 2.

Configure SVTI-2 as the static VTI of spoke 1 and SVTI-2 as the static VTI of spoke 2 (refer the above illustration). Tunnel source for SVTI-2 should be the same outside interface.

2. For each spoke, configure the routing policy. For more information, see Configure routing for multiple hubs in a route-based VPN, on page 1763.
3. Verify the configuration and tunnel statuses. For more information, see Verify multiple hubs configuration in a route-based VPN, on page 1764.

Configure routing for multiple hubs in a route-based VPN

Configure dynamic routing using BGP on both hub and spoke devices, and implement Policy Based Routing on spokes to enable proper traffic routing in a route-based VPN topology with multiple hubs.

Before you begin

Configure topology 1 and 2 as explained in Configure multiple hubs in a route-based VPN, on page 1761.

Procedure

Step 1 Configure dynamic routing for the hub using BGP.

a) Choose **Devices** > **Device Management**, and click **Edit** ( ) then click the **Routing** tab.
b) In the left pane, choose **General Settings > BGP**.
c) Check the **Enable BGP** check box and enter the **AS number**.

You can configure the other fields AS per your requirement.

d) Click **Save**.
e) On the left pane, choose **BGP > IPv4**.
f) Check the **Enable IPv4** check box.
g) Click the **Neighbor** tab, click **Add** and configure these parameters.

1. **IP Address**: Enter the tunnel interface IP address of Spoke 1.
2. **Remote AS**: AS number of Spoke 1.
3. Check the **Enabled Address** check box.
4. Click **OK**.

Repeat the above steps to add Spoke 2 as a neighbor.

h) Click **Save**.
i) Click the **Networks** tab and click **Add** to advertise the network behind the hub to the peers.

Step 2 Configure dynamic routing for the spokes using BGP.

The BGP configuration for the spokes is similar to that of the hub except for these differences:

- Configure Hub 1 and Hub 2 AS the neighbors for both the spokes and use the tunnel interface IP address of the hubs.
- When you configure networks, use the network behind each spoke.

Step 3 Configure Policy Based Routing on the spokes.

a) In the left pane, choose **Policy Based Routing** and click **Add**.
b) Choose the **Ingress Interface** from the drop-down list.
c) Click **Add** to configure a Match ACL.

For example, for spoke 1, source network is 192.168.20.0/24 and destination network is 192.168.10.0/24.

d) Choose Egress Interfaces from the **Send to** drop-down list.
e) Choose Order from the **Interface Ordering** drop-down list.
f) Select the SVTI-1 and SVTI-2 interfaces AS the egress interfaces.
g) Click **Save**.

If you want to use the hubs as a load-balancing pair, you must configure ECMP.

Step 4 Deploy the configurations on the hub and spokes.

What to do next

Verify the configurations and tunnel statuses. For more information, see Verify multiple hubs configuration in a route-based VPN, on page 1764.

Verify multiple hubs configuration in a route-based VPN

To verify the multiple hub configurations and the tunnel statuses.

- After the deployment, check the tunnel status.
- Use these show commands for each endpoint to verify the configurations:
  - **show run route-map**
  - **show run access-list**
  - **show route-map**
  - **show route**

Route traffic through a backup VTI tunnel

A backup VTI tunnel ensures traffic continuity by providing an alternative path when the primary VTI tunnel fails to route traffic in route-based VPN configurations.

You can deploy the backup VTI tunnel in these scenarios:

- Both peers having service provider redundancy backup.

  In this case, there are two physical interfaces, acting as the tunnel sources for the two VTIs of the peers.

- Only one of the peers has service provider redundancy backup.

  In this case, there's an interface backup on only one side of the peer and on the other end, there is only one tunnel source interface.

Before you begin

Ensure that you review the guidelines and limitations. For more information, refer to Guidelines for virtual tunnel interfaces, on page 1749.

Procedure

Step 1 Create the VTI interface.

For more information, refer to Add a VTI interface, on page 1752

Step 2 In the **Add Endpoint** dialog box of the **Create New VPN Topology** wizard, click **Add Backup VTI** to configure the respective backup interface for each peer.

For more information, refer to

- Configure endpoints for a point-to-point VPN topology, on page 1755
- Configure endpoints for a hub and spoke topology, on page 1757

Step 3 Choose **Devices > Device Management**, and edit the threat defense device and click **Routing** to configure the routing policy.

Step 4 Choose **Policies > Access Control** to configure the access control policy.

Configure dynamic VTI for a route-based site-to-site VPN

To configure dynamic VTI for a route-based site-to-site VPN in the Management Center:

Step	Do This	More Information
1	Create a dynamic VTI interface on the hub.	Add a VTI interface, on page 1752
2	Create static VTI interfaces on the spokes.	Add a VTI interface, on page 1752
3	Create a route-based site-to-site VPN.	Create a route-based site-to-site VPN, on page 1754
4	Configure the routing policy and the access control policy.	Configure endpoints for a hub and spoke topology, on page 1757

Configure a virtual router with dynamic VTI

To configure a virtual router with dynamic VTI for a route-based site-to-site VPN in the Management Center:

Step	Do This	More Information
1	Create a route-based site-to-site VPN with a dynamic VTI interface on the hub and static VTI(s) on the spoke(s).	Create a route-based site-to-site VPN, on page 1754
2	Create a virtual router.	Create a virtual router, on page 1386

Step	Do This	More Information
3	Assign the interfaces to the virtual router.	Configure a virtual router, on page 1387
4	Configure routing policies for the hub and spoke.	Configure endpoints for a hub and spoke topology, on page 1757
5	Configure access control policies for the hub and spoke.	Configure endpoints for a hub and spoke topology, on page 1757

Configure routing and AC policies for VTI

After you configure VTI interfaces and the VTI tunnel on both the devices, you must configure these parameters:

- A routing policy to route VTI traffic between the devices over the VTI tunnel.
- An access control rule to allow encrypted traffic.

Procedure

Step 1 Configure routing for VTI interfaces.

For the VTI interfaces, you can configure static route or routing protocols such AS BGP, EIGRP, OSPF/OSPFv3.

a. Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b. Click **Routing**.

c. Configure static route, or BGP, EIGRP, OSPF/OSPFv3.

Routing	Parameters	More Information
Static Route	• **Interface**—Select the VTI interface. For a backup tunnel, select the backup VTI interface. • **Selected Network**—Remote peer's protected network. • **Gateway**—Remote peer's tunnel interface IP address. For a backup tunnel, select the remote peer's backup tunnel interface IP address. • **Metric**—For a backup tunnel, configure a different metric to handle the failover of the traffic flow over the backup tunnel. **Note** DVTI does not support static routes.	Add a static route, on page 1361

Routing	Parameters	More Information

BGP	• Under **General Settings** > **BGP**, enable BGP, provide the AS number of the local device, and add Router ID (if you choose Manual).
	• Under **BGP**, enable IPv4/IPv6 and click the **Neighbor** tab to configure the neighbors.
	  • **IP Address**—Remote peer's VTI interface IP address. For a backup tunnel, add a neighbor with the remote peer's backup VTI interface IP address.
	  • **Remote AS**—Remote peer's AS number.
	• Click the **Redistribution** tab, select the **Source Protocol** AS Connected to enable connected route redistribution.
		Configure BGP, on page 1496

Routing	Parameters	More Information
EIGRP	• Enable EIGRP, provide the AS number of the local device, and select the networks or hosts that participate in the EIGRP routing process. • Click the **Neighbors** tab and define the static neighbors for the EIGRP process. • To advertise summary addresses from a VTI interface, click the **Summary Address** tab, choose the VTI interface from the **Interface** drop-down. From the **Network** drop-down, choose the network to be summarized. • Click the **Interfaces** tab to configure the interface-specific EIGRP routing properties for the VTI interface. To enable EIGRP split-horizon on the interface, check the **Split Horizon** check box. You can also configure the **Hold Time** that is advertised by the device in the EIGRP hello packets.	
OSPF	• Check the **Process 1** check box, and choose the OSPF role. • Click the **Interface** tab and choose a VTI interface.	OSPFv2 configuration, on page 1454
OSPFv3	• Check the **Process 1** and **Enable Process 1** check boxes, and choose the OSPFv3 role. • Click the **Interface** tab and choose a VTI interface.	OSPFv3 configuration, on page 1467

Step 2 Add an access control rule to the access control policy on the device to allow encrypted traffic between the VTI tunnels with these settings:

a. Create the rule with the Allow action.

b. Select the VTI security zone of the local device AS the source zone and the VTI security zone of the remote peer AS the destination zone.

c. Select the VTI security zone of the remote peer AS the source zone and the VTI security zone of the local device AS the destination zone.

For more information about configuring an access control rule, see Create and edit access control rules, on page 2066.

View virtual tunnel interface details

You can view details of dynamic and static VTIs of route-based VPNs on the device. For every VPN topology, you can also view details of all the dynamically generated virtual access interfaces associated with each dynamic VTI.

Before you begin

- Supported device for static VTIs: Firewall Threat Defense Version 7.0 and later
- Supported device for dynamic VTIs: Firewall Threat Defense Version 7.3 and later

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click the **Virtual Tunnels** tab.

For each VTI, you can view details such as name, IP address, IPsec mode, tunnel source interface details, topology, and remote peer IP. You can also view if path monitoring is enabled for each interface.

Deploy a SASE tunnel on Umbrella

Guidelines for configuring SASE tunnels on Umbrella

SASE topology supports:

- Only PSK-based authentication
- IKEv2
- High availability

General configuration guidelines

- Cloud-Delivered Firewall Management Center does not discover tunnels created directly on Umbrella or by other applications.

- You can add only devices managed by the Cloud-Delivered Firewall Management Center as endpoints for the SASE topology. You cannot add extranet devices.

  For high availability pairs, the HA pair names appear in the endpoint list.

- If you check the **Deploy configuration on threat defense nodes** check box in the wizard, the Umbrella SASE topology configuration is deployed on the Firewall Threat Defense only after the tunnels are deployed on Umbrella.

  The Cloud-Delivered Firewall Management Center requires the local tunnel ID to deploy the Umbrella configuration in Firewall Threat Defense. Umbrella generates the complete tunnel ID (*<prefix>@<umbrella generated ID>*-umbrella.com) only after the Cloud-Delivered Firewall Management Center deploys the tunnel on Umbrella.

Version migration guideline (Pre-version 7.3)

You must create new SASE topologies in version 7.3 and delete the existing SASE topologies because Cloud-Delivered Firewall Management Center does not recognize topologies using Umbrella data center as an extranet hub created before version 7.3 as SASE topologies.

Deployment status and monitoring guideline

You can track the tunnel deployment status in three places:

- Cisco Umbrella Configuration dialog box of the wizard
- Notifications page under the Deployments and Tasks tabs
- Site to Site VPN Monitoring dashboard

Tunnel deletion guidelines

- You cannot edit or delete a SASE topology while deployment to Umbrella is in progress.
- When you delete a tunnel from Cloud-Delivered Firewall Management Center and if it is unable to delete the tunnel from Umbrella, you must manually delete it by logging into Umbrella.

Limitations for configuring SASE tunnels on Umbrella

SASE topology does not support these features:

- Clustering
- Certificate-based authentication
- IKEv1

Deploy a SASE tunnel on Umbrella

This section provides instructions to deploy a SASE tunnel on Umbrella from a Firewall Threat Defense device using the Cloud-Delivered Firewall Management Center.

Step	Do This	More Info
1	Review the guidelines and limitations.	Guidelines for configuring SASE tunnels on Umbrella, on page 1769

Step	Do This	More Info
2	Ensure that you meet the prerequisites.	Prerequisites for configuring Umbrella SASE tunnels, on page 1771
3	Configure Cisco Umbrella connection settings.	• Configure Cisco Umbrella Connection Settings • Map Management Center Umbrella parameters and Cisco Umbrella API keys, on page 1771
4	Configure a SASE tunnel on Umbrella.	Configure a SASE tunnel for Umbrella, on page 1773
5	View the status of the SASE tunnel.	View SASE tunnel status, on page 1775

Prerequisites for configuring Umbrella SASE tunnels

License prerequisites

- You must have a Cisco Umbrella Secure Internet Gateway (SIG) Essentials subscription.
- You must enable your Smart License account with the export-controlled features to deploy tunnels on Umbrella from the Cloud-Delivered Firewall Management Center. If this license is not enabled, you can only create a SASE topology. You cannot deploy tunnels on Umbrella.

General prerequisites

- You must establish an account with Cisco Umbrella at https://umbrella.cisco.com, log into Umbrella at http://login.umbrella.com, and obtain the required information to establish connection to Cisco Umbrella.
- You must register Cisco Umbrella with the Cloud-Delivered Firewall Management Center and configure the management key and the management secret in the Cisco Umbrella Connection Settings. The management center requires the management key and management secret to fetch the datacenter details from the Cisco Umbrella cloud. You must also configure the Organization ID, Network Device Key, Network Device Secret, and the Legacy Network Device Token in the Cisco Umbrella Connection Settings.

  For more information, refer to:

  - Configure Cisco Umbrella Connection Settings
  - Map Management Center Umbrella parameters and Cisco Umbrella API keys, on page 1771

- Ensure that Umbrella data center is reachable from the Firewall Threat Defense.
- You can deploy a tunnel only between Cisco Umbrella and a Firewall Threat Defense version 7.1.0 and later.

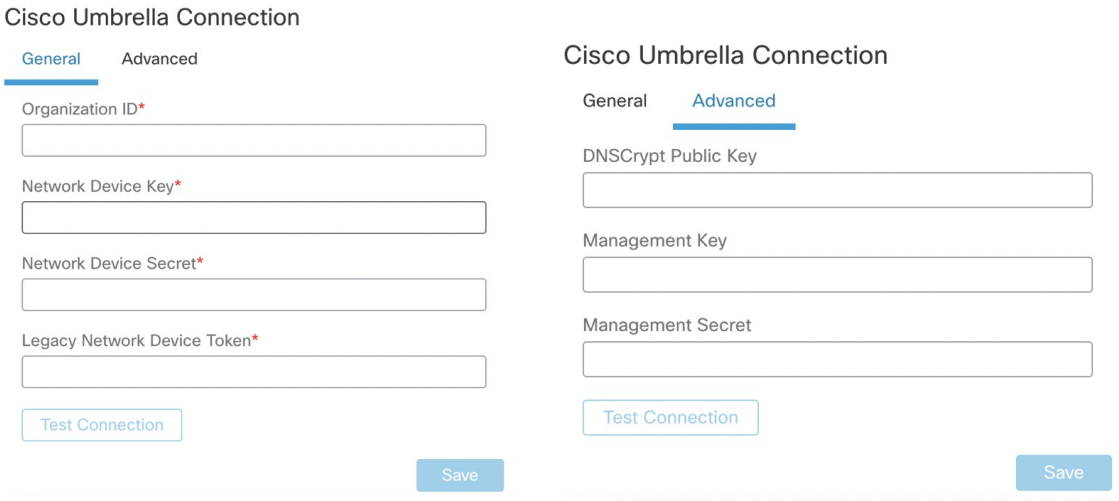
Map Management Center Umbrella parameters and Cisco Umbrella API keys

To register Cisco Umbrella with the Cloud-Delivered Firewall Management Center and configure the Umbrella parameters in the Cloud-Delivered Firewall Management Center, you must complete these steps:

1. Log in to Cisco Umbrella.
2. Choose **Admin > API Keys > Legacy Keys**.

3. Generate and copy the required API keys.
4. Use the API keys to configure the Cisco Umbrella connection parameters in the Cloud-Delivered Firewall Management Center.

This figures show the parameters that you must configure in the Cisco Umbrella Connection in the Cloud-Delivered Firewall Management Center. DNScrypt Public Key is an optional parameter.



This figure shows the Cisco Umbrella API keys that you must use to register the Cisco Umbrella with the Cloud-Delivered Firewall Management Center.

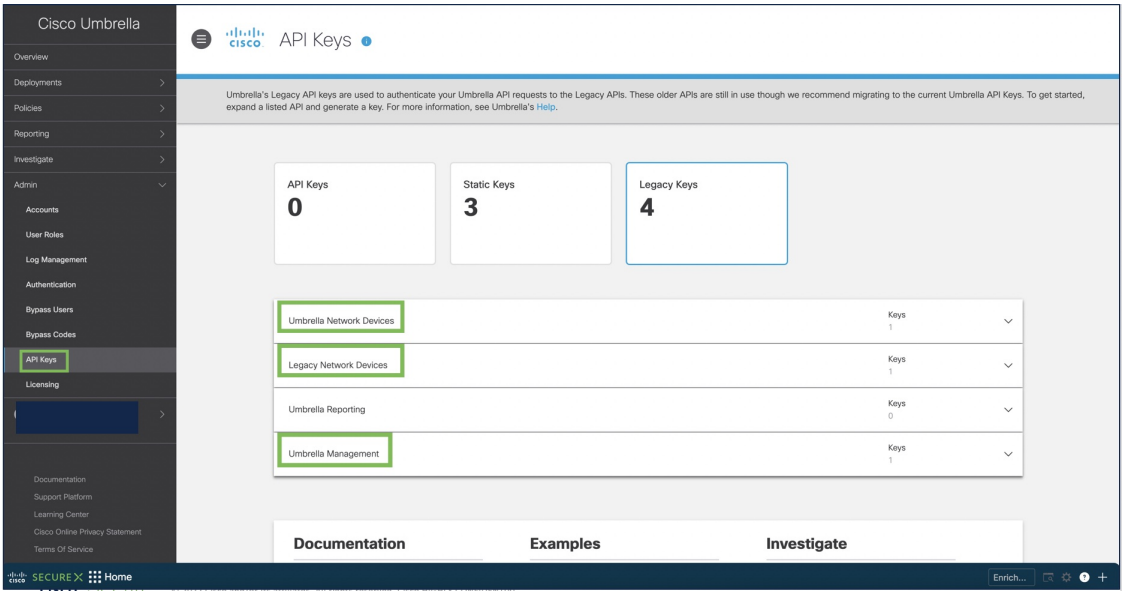


Table 101: Map Cloud-Delivered Firewall Management Center Umbrella Parameters and Cisco Umbrella API Keys

Management Center Parameter	Cisco Umbrella API Keys
Network Device Key Network Device Secret	Umbrella Network Devices
Legacy Network Device Token	Legacy Network Devices
Management Key Management Secret	Umbrella Management

Configure a SASE tunnel for Umbrella

Configure a SASE tunnel to establish secure connectivity between threat defense devices and Cisco Umbrella data centers for cloud-delivered security services.

SASE (Secure Access Service Edge) tunnels provide secure connections to Umbrella data centers, enabling cloud-based security functions. These tunnels use IKEv2 with pre-shared keys for authentication and static VTI interfaces for routing traffic.

Before you begin

Ensure that you review the prerequisites and guidelines in Prerequisites for configuring Umbrella SASE tunnels, on page 1771 and Guidelines for configuring SASE tunnels on Umbrella, on page 1769.

Follow these steps to configure a SASE tunnel for Umbrella:

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**, click **+SASE Topology** and enter a unique **Topology Name**.

Step 2 **Pre-shared Key**: This key is auto-generated according to the Umbrella PSK requirements. For a single topology, the pre-shared key is common for all threat defense spokes and Umbrella.

The device and Umbrella share this secret key, and IKEv2 uses it for authentication. If you want to configure this key, it must be between 16 and 64 characters in length, include at least one uppercase letter, one lowercase letter, one numeral, and have no special characters. Each topology must have a unique pre-shared key. If a topology has multiple tunnels, all the tunnels have the same pre-shared key.

Step 3 Choose a data center from the **Umbrella Data center** drop-down list. (Configure routing on the Firewall Threat Defense to ensure reachability of the umbrella DC from the Firewall Threat Defense.)

Step 4 Click **Add** to add a Firewall Threat Defense node.

a) Choose a Firewall Threat Defense from the **Device** drop-down list.

Only devices managed by the Cloud-Delivered Firewall Management Center appear in the list. For high availability pairs, the HA pair names appear in the endpoint list.

b) Choose a static VTI interface from the **VPN Interface** drop-down list.

To create a new static VTI interface, click Add +. The **Add Virtual Tunnel Interface** dialog box appears with the following pre-populated default configurations.

- Tunnel Type is static.

- Name is <*tunnel_source interface logical name*>+ static_vti +<*tunnel ID*>. For example, outside_static_vti_2.
- Tunnel ID is auto-populated with a unique ID.
- Tunnel Source Interface is auto-populated with an interface with an 'outside' prefix.
- IPsec tunnel mode is IPv4.
- IP address is from the 169.254.x.x/30 private IP address range.

c) Enter a prefix for the local tunnel ID in the **Local Tunnel ID** field.

The prefix can have a minimum of eight characters and a maximum of 100 characters. Umbrella generates the complete tunnel ID (<*prefix*>@<*umbrella generated ID*>-umbrella.com) after the management center deploys the tunnel on Umbrella. The management center then retrieves and updates the complete tunnel ID and deploys it on the threat defense device. Each tunnel has a unique local tunnel ID.

d) Click **Save** to add the endpoint device to the topology.

You can add multiple endpoints in a SASE topology.

Step 5 Click **Next** to view the summary of the Umbrella SASE tunnel configuration.

- **Endpoints** pane: Displays the summary of the configured endpoints.
- **Encryption Settings** pane: Displays the default IKEv2 policies and IKEv2 IPsec transform sets for the topology.

Step 6 Check the **Deploy configuration on threat defense nodes** check box to trigger deployment of the network tunnels to the threat defense. This deployment occurs after the tunnels are deployed on Umbrella. Local tunnel ID is required for the threat defense deployment.

Step 7 Click **Save**.

This action:

a. Saves the topology in the management center.

b. Triggers deployment of the network tunnels to the Umbrella.

c. Triggers deployment of the network tunnels to the threat defense devices, if the option is enabled. This action commits and deploys all the updated configurations and policies, including non-VPN policies, since the last deployment on the device.

d. Opens the **Cisco Umbrella Configuration** window and displays the status of the tunnel deployment on Umbrella. For more details, see View SASE tunnel status, on page 1775.

What to do next

For the interesting traffic intended to flow through the SASE tunnel, configure a PBR policy with a specific match criteria to send the traffic through the VTI interface.

Ensure that you configure a PBR policy for each endpoint of the SASE topology.

View SASE tunnel status

After configuring Umbrella SASE topology, you need to monitor the tunnel deployment status and verify that the configuration has been successfully pushed to Umbrella.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**.

Step 2 Click **+SASE Topology**.

Step 3 Enter a unique **Topology Name**, **Pre-shared Key**, choose a data center, add a device, and click **Next**.

Step 4 View the summary of the Umbrella SASE Tunnel configuration and click **Save**.

The **Cisco Umbrella Configuration** window appears.

You can view the topology details such as name, data center, data center IP address, start, and end time of the tunnel deployment.

You can view the deployment status of the tunnels on Umbrella. The different tunnel deployment statuses are:

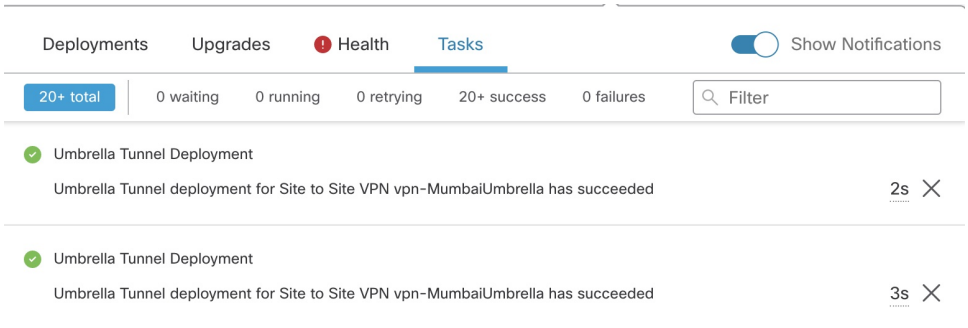
- Pending: The management center hasn't pushed the configuration to Umbrella.
- Success: The management center successfully configured a tunnel on Umbrella.
- In Progress: The management center is deploying the tunnel on Umbrella.
- Failure: The management center couldn't configure a tunnel on Umbrella.

If the status appears as pending, or failure, use the transcript to troubleshoot the tunnel creation. Click the **Transcript** button to view the transcript details such as the APIs, request payload, and the response received from Umbrella.

Step 5 Click **Umbrella Dashboard** to view the network tunnels in Cisco Umbrella.

Step 6 View the Umbrella tunnel deployment status in:

- The **Notifications** page under the **Deployments** and **Tasks** tabs.



Secure traffic using Cisco Secure Access and Firewall Threat Defense devices

Cisco Secure Access is a cloud-based security service edge (SSE) solution that protects against internet threats and enables secure access to the internet, SaaS applications, and private resources from any location.

Benefits of using Secure Access

Secure Firewall seamlessly integrates with Secure Access for various use cases primarily for:

- **Enhanced visibility and control for remote users**: Secure Access extends security policies beyond the physical branch perimeter to protect users regardless of location, providing the same level of threat protection and access controls equivalent to on-premises, creating consistent security enforcement.
- **Simplified security architecture**: Secure Access integrates with a Firewall Threat Defense device to consolidate multiple security functions such as Secure Web Gateway, Cloud Access Security Broker (CASB), Data Loss Prevention (DLP), and cloud firewall. This consolidation provides a single-pane-of-glass experience for policy management across branch deployments and cloud-delivered security services.

Benefits of using SASE wizard to automate Secure Access tunnels from Firewall Threat Defense

You can automatically create secure tunnels between Firewall Threat Defense and Secure Access using Cloud-Delivered Firewall Management Center. The SASE wizard provides a guided workflow to handle the tunnel setup across both the products, helping you avoid switching between consoles or manually matching tunnel parameters.

- Reduce tunnel setup time and configuration errors.
- Secure internet, cloud, and private application access through Secure Access.
- Simplify operations for distributed enterprise environments.

Licenses for automatic tunnel between Secure Access and Firewall Threat Defense devices

This topic lists the required licenses for using the SASE wizard to establish automatic tunnels between Threat Defense devices and Secure Access.

- Cloud-Delivered Firewall Management Center Essentials (formerly Base) license must allow export-controlled functionality. Choose **System** (⚙) > **Licenses** > **Smart Licenses** to verify this in Cloud-Delivered Firewall Management Center.

Note Evaluation license does not support this feature.

- Cisco Secure Access Essentials subscription license must be available.

Prerequisites for an automatic tunnel between Secure Access and Firewall Threat Defense devices

This topic provides the prerequisites required to establish an automatic tunnel between Secure Access and Firewall Threat Defense devices for secure traffic integration.

- Integrate a Cisco Security Cloud Control tenant with a Secure Access organization.
- Configure basic routing and access control policies to route traffic between the Secure Access data center and the Firewall Threat Defense device.
- Firewall Threat Defense devices must be Version 7.2.0 and later.

Guidelines for an automatic tunnel between Secure Access and Firewall Threat Defense devices

General guidelines

- Set up routing in the Firewall Threat Defense device to route traffic over the tunnel.
- Delete the Firewall Threat Defense device from the topology and add it again with the required routing configuration to update routing settings in a SASE topology.
- When change management is enabled, tunnels are not deployed automatically to Secure Access. Manually deploy tunnels to Secure Access from the **Site-to-Site VPN** page after the change management ticket is approved.

Multi-ISP SASE topology guidelines

- To create a multi-ISP SASE topology, you must establish multiple tunnels from a Firewall Threat Defense device to the same Secure Access region using different interfaces. Since each SASE topology supports only one WAN interface, use the SASE wizard to create multiple topologies.
- When you configure multi-ISP SASE topologies that use the same region and device but different VPN interfaces, ensure that all topologies have identical settings. The tunnel ID prefix, passphrase, and routing configurations must match across all topologies.
- To update routing settings in a multi-ISP SASE topology, delete the Firewall Threat Defense device from each topology and add it again with the required configuration. Ensure that routing settings are consistent across all topologies.

Limitation for automatic tunnel between Secure Access and Firewall Threat Defense devices

Clustering is not supported.

Configure automatic tunnel between Secure Access and Firewall Threat Defense devices using SASE wizard

The SASE wizard simplifies tunnel creation from Firewall Threat Defense devices to Secure Access by automating multiple steps.

Before you begin

Ensure that you review Prerequisites for an automatic tunnel between Secure Access and Firewall Threat Defense devices, on page 1777 .

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site** , and click **Add** .

Step 2 In the **Topology Name** field, enter a name for the SASE topology.

Step 3 Click the **SASE Topology** radio button and click **Create** .

Step 4 Configure a Secure Access region by choosing a region from the **Region** drop-down list:

A Secure Access region is a cluster of data centers in a specific geographic area.

Step 5 Configure Firewall Threat Defense nodes:

a) Click **Add** and configure the parameters in the **Add Threat Defense Node** dialog box.

b) From the **Device** drop-down list, choose a Firewall Threat Defense device.

Note
Only Firewall Threat Defense devices managed by Cloud-Delivered Firewall Management Center are displayed in this list. Extranet devices won't appear in the list.

c) From the **VPN Interface** drop-down list, choose a WAN-facing or internet-facing physical interface that establishes a VPN connection with Secure Access .

Step 6 Configure **Tunnel ID and passphrase** .

a) In the **Prefix for tunnel ID** field, enter a prefix containing 8 to 30 characters. The prefix can contain alphabets, numbers, period (.), underscore (_), and dash (-).

Secure Access uses this prefix to generate the complete tunnel ID after the Cloud-Delivered Firewall Management Center deploys the tunnel in Secure Access . The tunnel ID format is < *prefix* >@< *org* >< *hub* >.sse.cisco.com.

b) The passphrase is auto generated. However, you can regenerate or enter a different passphrase. In the **Passphrase** field, enter a passphrase containing 16 to 64 characters. The passphrase must contain at least one upper-case alphabet, one lower-case alphabet, and one number, and cannot include special characters.

c) Confirm the passphrase.

d) Click **Next** .

Step 7 Configure **NAT or routing** :

a) Check the **Enable NAT/outbound only** check box to enable Network Address Translation (NAT) for correct routing and identification.

Use this option if the IP address of the subnet behind this tunnel group overlaps with the IP address in your network. When this option is selected, the tunnel group supports only outbound traffic and cannot provide access to private applications hosted at this site.

Note
Routing options are unavailable when the network contains overlapping subnets.

b) Configure one of these options for **Routing** :

- Click the **Static routing** radio button.

  Use the **Network IP addresses** field to add the IP addresses of networks for this tunnel group. Add all the public and private addresses used internally by your organization.

- Click the **Dynamic routing** radio button.

  Use this option to advertise your internal networks over BGP.

  In the **AS number** field, enter the AS number of the device. This value must match the device's BGP AS number in the routing settings of the device.

c) Click **Next** .

Step 8 Check the **Deploy to Threat Defense devices** check box to trigger deployment of all the configurations besides the Secure Access auto tunnel configurations that are yet to be deployed on the device.

Step 9 Click **Finish** to save and validate the configurations, and create the SASE topology.

The wizard performs these actions:

- Saves all Firewall Threat Defense device configurations in Cloud-Delivered Firewall Management Center .
- Pushes the configurations to Secure Access.
- Creates virtual tunnel interfaces (VTIs) on the device.
- Creates a SASE VPN topology.
- Triggers deployment of all the configurations besides the Secure Access auto tunnel configurations that are yet to be deployed in the device if the **Deploy to Threat Defense devices** option is enabled.
- Opens the **Configuration of Secure Access Tunnels** dialog box, which displays the status of the tunnel deployment in Secure Access .

  Click the **Transcript Details** ( ) button to view transcript details such as the APIs, request payload, and the responses received from Secure Access .

You can view the SASE topology in the **Site-to-Site VPN & SD-WAN** page (**Devices** > **VPN** > **Site to Site**).

What to do next

1. Create an extended Access Control List (ACL).

   This ACL defines the specific DNS and web traffic intended for routing through the tunnel to Secure Access . For more information, refer to Configure an extended ACL object, on page 2949 .

2. Configure the static route to the next hop of the VTI interface.
3. Create a policy-based routing (PBR) policy.

   Use the newly created extended ACL within a policy-based routing policy to direct the defined DNS and web traffic through the tunnel to Secure Access for security inspection. For more information, refer to Configure policy-based routing policy, on page 1548 .
4. When you create multiple SASE topologies for a multi-ISP setup, configure ECMP zones with the VPN interfaces to balance the load of application traffic.
5. Perform validation. For more information, refer to Validate Secure Access integration with Firewall Threat Defense devices, on page 1780 .

Validate Secure Access integration with Firewall Threat Defense devices

You can validate the integration between Secure Access and Threat Defense devices by checking tunnel statuses and viewing transcript details.

Verify tunnel statuses in Cloud-Delivered Firewall Management Center

To view the SASE topology tunnel details in the site-to-site VPN dashboard, choose **Overview** > **Dashboards** > **Site to Site VPN**.

To view more details about each tunnel:

1. For each tunnel, hover your cursor over a topology and click the **View** () icon.
2. Click the **CLI Details** tab.
3. Click **Maximize View**. You can view the output of these commands:
   - **show crypto ipsec sa peer**: Shows the number of packets that are transmitted through the tunnel.
   - **show VPN-sessiondb detail l2l filter ipaddress**: Shows detailed data for the VPN connection.

Verify tunnel statuses in Secure Access

1. In Secure Access, choose **Connect > Network Connections > Network Tunnel Groups**.
2. Click a network tunnel group to view details.

   In the **Network Tunnels** area, the primary and secondary tunnels will be displayed:
   - **Primary1** tunnel originates from the branch's outside interface and connects to the primary Secure Access data center.
   - **Secondary1** tunnel also originates from the branch's outside interface but connects to the secondary data center, serving as a standby for continuous connectivity.

View transcript details

In the **Configuration of Secure Access Tunnels** dialog box, click the **Transcript Details** () button to view the transcript details such as the APIs, request payload, and the responses received from Secure Access.

Transcript details are not available in these scenarios:

- When you delete a Network Tunnel Group (NTG) or an endpoint.
- When you configure any multi-ISP tunnel after configuring the primary tunnel.

Monitor site-to-site topologies

This section provides information about the different ways to monitor your site-to-site topologies.

Monitor site-to-site VPNs using Site-to-Site VPN summary page

You can view a summary of your site-to-site VPN topologies in the **Site-to-Site VPN Summary** page. For each topology, you can view details such as VPN type, network topology, VPN interfaces, VPN devices, and tunnel statuses.

You can edit or delete the topology using the edit and delete buttons. For SASE topology VPNs, you have options to deploy, edit and delete any topology.

Monitor site-to-site VPNs using Site-to-Site VPN dashboard

Choose **Overview** > **Dashboards** > **Site to Site VPN** to open the site to site dashboard.

The **Site-to-Site VPN** dashboard provides a snapshot and status of the site-to-site VPN tunnels. You can view the list of tunnels between peer devices and the status of each tunnel: Active, Inactive, or No Active Data. You can filter the data in the table according to the topology, device, and status. The dashboard shows live data and you can also configure the refresh interval.

Benefits of the dashboard

The dashboard provides these benefits:

- Identify problematic VPN tunnels and troubleshoot them.
- Verify connectivity between the site-to-site VPN peers devices.
- Monitor the health of the VPN tunnels to provide uninterrupted VPN connectivity between sites.
- Use Packet tracer to troubleshoot VPN tunnels.

Widgets of the dashboard

- **Tunnel Status**—A table listing the tunnel status of the site-to-site VPNs , including the SASE tunnels for Umbrella, configured using the Cloud-Delivered Firewall Management Center. For more information, refer to Tunnel status widget, on page 1782.
- **Tunnel Summary**—Aggregated status of the tunnels in a donut graph.
- **Topology**—Status of tunnels summarized by topology.

Status of VPN tunnels

The different VPN tunnel statuses that you can view in the Site-to-Site dashboard are:

- **Inactive**—A policy-based (crypto map-based) VPN tunnel is inactive if all the IPSec tunnels are down. A VTI or and SASE topology VPN tunnel is down if the tunnel encounters any configuration or connectivity issues.
- **Active**—In the Cloud-Delivered Firewall Management Center, policy-based site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. A policy-based VPN tunnel is in the Active state if the Cloud-Delivered Firewall Management Center identifies interesting traffic through the tunnel after the deployment. An IKE tunnel is up only if a minimum of one IPsec tunnel is up.

  Route-based VPN (VTI) and SASE topology VPN tunnels do not require interesting traffic to be in the Active state. They are in the Active state if they are configured and deployed without errors.
- **No Active Data**—Policy-based and SASE topology VPN tunnels remain in the No Active Data state until there is a traffic flow event through the tunnel for the first time. The No Active Data state also lists the policy-based and route-based VPNs that have been deployed with errors.

Important notes about tunnel statuses in Cloud-Delivered Firewall Management Center

- The VPN statuses in the Cloud-Delivered Firewall Management Center are event-based. The Cloud-Delivered Firewall Management Center does not initiate status updates. Hence, there might be mismatches between the tunnel statuses in the dashboard and the Firewall Threat Defense . You can view the correct status in the **CLI Details** tab of the **Tunnel Status** widget.
- When a Firewall Threat Defense switches over to a secondary Firewall Threat Defense , there is a mismatch between the status of the VPN tunnels in the Cloud-Delivered Firewall Management Center and the Firewall Threat Defense . When the device switches back to the primary device, the correct tunnel status appears.
- The Cloud-Delivered Firewall Management Center does not update the tunnel status of Firewall Threat Defense devices earlier than 7.3 after the devices reboot. We recommend that you bring the tunnel down using the command **vpn-sessiondb logoff index** and bring it up using the packet tracer.

Limitations of the dashboard

- The table shows the list of site-to-site , including SASE topology, VPNs that are deployed. It does not show the tunnels that are created and not deployed.
- The table does not show the information about the backup tunnels of policy-based VPNs and backup VTIs.
- For cluster deployments, the table does not show director change in real-time data. It shows only the director information that existed when the VPN was deployed. The director change reflects in the table only after the tunnel AM redeployed after the change.

Tunnel status widget

This widget lists the site-to-site VPNs, including SASE topology VPN, configured using the Cloud-Delivered Firewall Management Center and their tunnel statuses.

Features of the widget

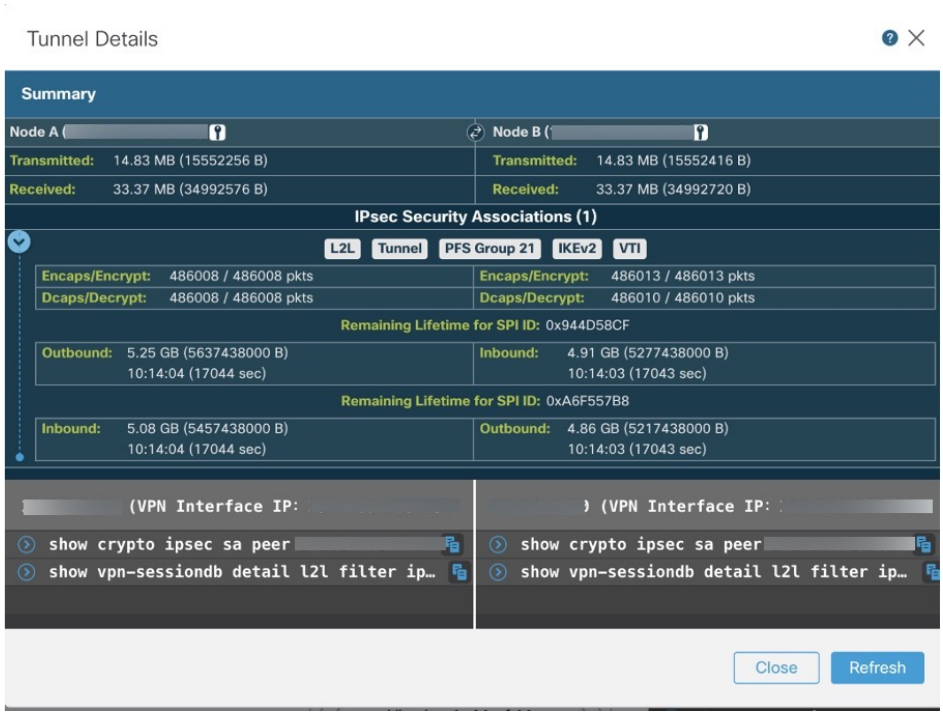
Hover over a topology and click **View** ( ) to view these details about the topology:

- **General**—Displays more information about the nodes such as IP address and interface name.
- **CLI Details**—Displays the CLI outputs for the following commands:
  - **show crypto ipsec sa peer**: Displays the IPsec SAs built between Node A and B.
  - **show vpn-sessiondb l2l filter ipaddress** : Displays information about VPN sessions.

  For an extranet device, no command output appears.

  Important details about the IKE and IPsec sessions, derived from the above command outputs, appear in a summarized, and user-friendly format. You can view the details of both nodes at a time. The icon next to the node name specifies the authentication type: preshared key or client certificate. The details include IKE statistics per tunnel and IPsec SA statistics as shown below:

Figure 369: Tunnel Status > View > CLI Details



- **Packet Tracer**—Use packet tracer to troubleshoot the VPN tunnels.

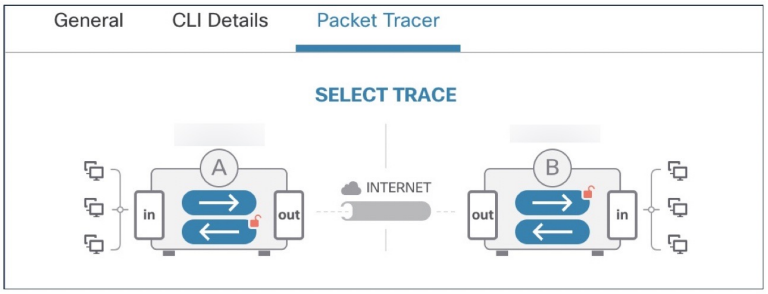
Packet Tracer

Packet tracer allows you to troubleshoot VPN tunnels between two threat defense devices. You can check if the VPN connection between device A and device B is up. This tool performs these actions:

- Injects a packet into the device and tracks the packet flow from the ingress to the egress ports.
- Simulates traffic after you configure the ingress interfaces of the devices along with the protected networks.

- Evaluates the packet against modules such as flow and route lookups, ACLs, protocol inspection, NAT, and QoS.

Figure 370: Packet Tracer



For each device, the tool runs an encrypted trace and a decrypted trace (packet is treated as decrypted VPN traffic). You can run four different traces between the ingress and egress ports of the devices. Click the individual encrypt and decrypt options to enable or disable the trace.

When you run the trace, the tool executes the trace sequentially in the following order:

1. Encrypted trace of A.
2. Decrypted trace of B.
3. Encrypted trace of B.
4. Decrypted trace of A.

After the trace completes, you can view the output of the trace with the results of each module.

Note You cannot run a decrypt trace for route-based (VTI-based) VPNs.

To run the Packet Tracer:

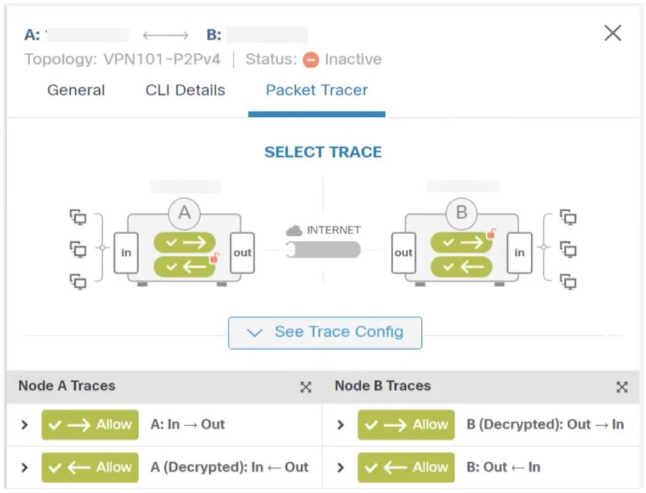
1. Click **See Detailed Config** to view the VPN interface name, VPN interface IP address, VTI interface name, and the VTI Interface IP address.
2. (Optional) Choose a protocol from the **Protocol** drop-down list. You can choose ICMP/8/0, TCP, or UDP.

   ICMP/8/0 is the default option. If you choose ICMP/8/0, 8 indicates the ICMP type as Echo Request and 0 indicates the ICMP code. If you choose TCP or UDP, choose the destination port from the **Destination Port** drop-down list. The range is from 0 to 65535.
3. Choose the ingress interface for both the devices on which to trace the packet from the **Ingress Interface** drop-down lists.
4. Enter an IP address from the same subnet as the ingress interface in the **Protected Network IP Address** fields.
5. Click **Trace Now**.

   After you initiate the trace, you can view if the trace is successful or not for each module. If the tunnel is down, the path appears in red. If the tunnel is up, the path appears in green. If a tunnel is down, click

Re-trace to run the tool again. For a crypto-map based VPN, when the tunnel is inactive with no interesting traffic, the initial trace can be red. Click **Re-trace** to run the trace again.

Figure 371: Packet Tracer after a Successful Trace



Extranet Nodes: You can initiate a packet trace for VPN tunnels with one node as an extranet. For an extranet node, you cannot choose the ingress interface. The remaining steps for the packet trace are the same. You can't run trace on the extranet side.

For example, if Node A is a managed threat defense and Node B is an extranet:

- Configure the ingress interface for node A.
- Configure the protected network for Node A and B.
- Click **Trace Now**. The traces appear for Node A and not for Node B.

Site-to-site VPN connection event monitoring

The site-to-site VPN connection event allows you to know if the VPN encrypts or do not encrypts the connection and helps you to troubleshoot connectivity issues, especially in multi-hop VPN deployments. The event dashboard of the Cloud-Delivered Firewall Management Center displays the IP address of the VPN peer (peer's IKE address) which encrypts or decrypts the traffic and displays the VPN action.

The VPN actions displayed are:

- If the connection is decrypted by the VPN, the column **Decrypt Peer** displays the IP address of the peer's address from where the flow was received and displays **Decrypt** as the VPN action.
- If the connection is encrypted by the VPN, the column **Encrypt Peer** displays the IP address of the VPN peer to which the flow is sent and displays **Encrypt** as the VPN action.
- If the VPN server cascades the connection, it gets decrypted on one tunnel and gets re-encrypted on another tunnel. In this case, both **Encrypt Peer** and **Decrypt Peer** IP addresses get appears in the event. The column **VPN Action** displays **VPN Routing** as the action to indicate that the connection transit through the VPN server.

If you enable the **Bypass Access Control Policy for decrypted traffic (sysopt permit-VPN)** option, the system bypasses the access control policy and do not log events for decrypted traffic. This option is disabled by default and all decrypted traffic in the VPN tunnel undergoes ACL inspection.

View site-to-site VPN connection events

Access the connection event viewer of the Cloud-Delivered Firewall Management Center to know if the VPN encrypts or do not encrypts the connection traffic and to retrieve the VPN peer details.

Before you begin

Ensure that you enable logging of connection events at the beginning of connection and at the end of connection in the access control rule.

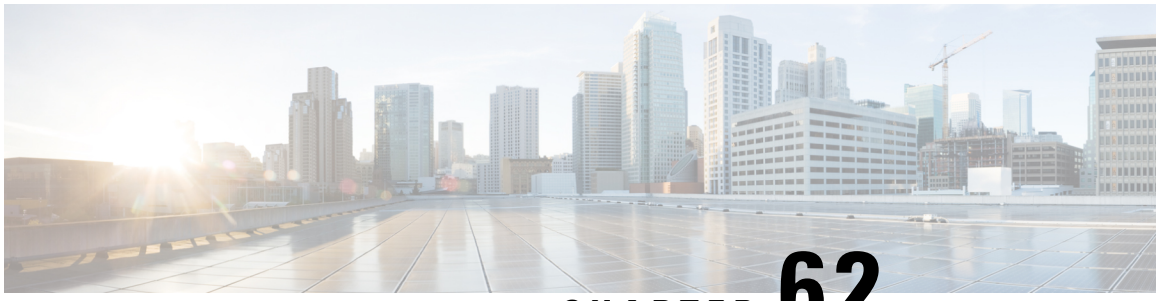
Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Choose these columns:

- Decrypt Peer
- Encrypt Peer
- VPN Action

Step 3 Click **Apply**.



CHAPTER 62

SD-WAN

This chapter describes the SD-WAN capabilities supported in Management Center.

- SD-WAN capabilities, on page 1787
- SD-WAN features, on page 1788
- SD-WAN wizard for secure branch network deployment, on page 1789
- Monitor SD-WAN topologies, on page 1805

SD-WAN capabilities

SD-WAN capabilities are software-defined networking features that

- replace traditional WAN routers and are agnostic to WAN transport technologies
- provide dynamic, policy-based, application path selection across multiple WAN connections, and
- support service chaining for additional services such as WAN optimization and firewalls.

SD-WAN capabilities in Secure Firewall

As organizations expand their operations across multiple branch locations, ensuring secure and streamlined connectivity becomes paramount. Deploying a secure branch network infrastructure involves complex configurations, which can be time-consuming and prone to configuration errors if not handled properly. However, organizations can overcome these challenges by leveraging Management Center and Threat Defense devices for a simplified and secure branch deployment.

By integrating Secure Firewall as a foundational component of the branch network architecture, organizations can establish a strong security baseline while simplifying the deployment process. This approach enables organizations to enforce unified security policies, optimize traffic routing, and ensure resilient connectivity.

Some of the SD-WAN capabilities supported on the Cisco Secure Firewall include:

- **Simplified management:**
  - SASE: Umbrella auto tunnel deployment
  - Dynamic VTI (DVTI) hub spoke topology simplification
- **Application awareness:**
  - Direct Internet Access (DIA) for public cloud and guest user

- Policy based routing (PBR) using applications as a match criteria
- Local tunnel ID support for Umbrella

- **Increased usable bandwidth:**
  - ECMP support for load balancing across multiple ISPs and VTIs
  - Application-based load balancing using PBR
- **High availability with near zero network downtime:**
  - Dual ISP configuration
  - Optimal path selection based on application-based interface monitoring.
- **Secure elastic connectivity:**
  - Route-based (VTI) VPN tunnels between headquarters (hub) and branches (spokes)
  - IPv4 and IPv6 BGP, IPv4 and IPv6 OSPF, and IPv4 EIGRP over VTI
  - DVTI hubs that support spokes with static or dynamic IP

SD-WAN features

This table lists the SD-WAN features of Management Center, including their introduction releases and links to detailed information.

Table 102: SD-WAN features

Feature	More Information
Secure traffic using Cisco Secure Access and Firewall Threat Defense devices	Secure traffic using Cisco Secure Access and Firewall Threat Defense devices, on page 1776
Configure settings to carry SGT over the SD-WAN tunnels using SD-WAN wizard.	SD-WAN wizard for secure branch network deployment, on page 1789
Support for ECMP in hubs and BFD for SVTI and DVTI.	SD-WAN wizard for secure branch network deployment, on page 1789
Use PBR to handle traffic based on advanced custom application detector.	Policy-Based Routing, on page 1541
Management of Remote Branch (HA) with Dual WAN Support	SD-WAN summary dashboard, on page 1805
PBR Support for Custom Applications	Policy-Based Routing, on page 1541
SD-WAN Wizard	SD-WAN wizard for secure branch network deployment, on page 1789

Feature	More Information
Application monitoring using SD-WAN Summary dashboard	SD-WAN summary dashboard, on page 1805
SD-WAN Summary Dashboard	SD-WAN summary dashboard, on page 1805
Policy-based routing with user identity and SGTs	Policy-Based Routing, on page 1541
Policy-based routing using HTTP path monitoring	Policy-Based Routing, on page 1541
Loopback interface support for VTIs	Loopback interfaces, on page 1052
Support for dynamic VTI (DVTI) with site-to-site VPN	Dynamic VTI, on page 1747
Umbrella auto tunnel	Deploy a SASE Tunnel on Umbrella
Support for IPv4 and IPv6 BGP, IPv4 and IPv6 OSPF, and IPv4 EIGRP for VTIs	BGP, on page 1491 OSPF, on page 1447 EIGRP, on page 1479
Route-based site-to-site VPN with hub and spoke topology	Create a route-based site-to-site VPN, on page 1754
Policy-based routing with path monitoring	Policy-Based Routing, on page 1541
Site to Site VPN Monitoring Dashboard	Monitor Site-to-Site VPNs Using Site-to-Site VPN Dashboard
Direct Internet Access/Policy Based Routing	Policy-Based Routing, on page 1541
Equal-Cost-Multi-Path (ECMP) zone with WAN interfaces	ECMP, on page 1431
ECMP zone with VTI interfaces	ECMP, on page 1431
Backup VTI for route-based site-to-site VPN	Route traffic through a backup VTI tunnel, on page 1764
Support for static VTI (SVTI) with site-to-site VPN	Static VTI, on page 1746

SD-WAN wizard for secure branch network deployment

The SD-WAN wizard in Management Center

- simplifies and automates VPN tunnel and routing configuration between centralized headquarters (hubs) and remote branch sites (spokes)
- creates route-based VPN tunnels and automates tasks such as generating tunnel interfaces, assigning IP addresses, and configuring BGP for the SD-WAN overlay network

- provides seamless routing by enabling hubs to act as route reflectors that determine the best routing path based on spokes' active and backup tunnels.

Hub and spoke network components

The SD-WAN wizard configures two types of network components:

- Hubs: Devices that enable secure VPN connectivity to and from one or more remote branch devices or spokes. Hubs also act as a gateway for spokes to communicate with each other.
- Spokes: Devices in remote branches that connect over VPN to a hub to securely access the corporate resources behind the hub. Spokes communicate with each other through the hub.

The SD-WAN wizard provides these benefits:

- Simplifies and automates the VPN and routing configuration of your SD-WAN network.
- Requires minimal user input
- Easily adds multiple branches at a time
- Provides easy dual ISP configurations
- Enables network scaling

Guidelines for using SD-WAN wizard

General guidelines

- When you create security zones or interface groups, choose **Routed** as the **Interface Type**.
- Use the spoke security zone to configure an access control policy that allows tunnel traffic to and from the spokes.
- If a spoke is part of multiple SD-WAN topologies, ensure that you use the same local community tag and learned route community tag in each SD-WAN topology. Note that the local community tags and learned route community tags must be different from each other.
- Configure unique local IKE identity for all tunnels across all your SD-WAN VPN topologies.
- Ensure that the spokes in an SD-WAN topology do not have the same protected network.

IP address configuration guidelines

- When you configure IP address pools for spokes, ensure the following:
  - The **Allow Overrides** check box must be unchecked.
  - If you are using multiple pools, the IP addresses of the pools must not overlap.
  - IP addresses must not overlap with any of the interfaces on the spoke.
- If a device has only IPv6 address configurations, you must configure the BGP router ID with a loopback or physical interface that has an IPv4 address (**Devices** > **Device Management**, and click **Routing** > **General Settings** > **BGP**).

Dual-ISP topology guidelines

- In a dual-hub SD-WAN topology, the hubs can be in different geographic locations and have different protected networks behind them. To ensure direct communication between these networks, ensure that you configure the following:
  - A point-to-point route-based VPN topology between the two hubs (**Devices** > **VPN** > **Site to Site**, and click **Add** > **Route-Based VPN**).
  - A dynamic routing protocol between the hubs (**Devices** > **Device Management**, click the device name and click **Routing**).
- In SD-WAN topologies with dual ISPs on spokes, the tunnel identity and the tunnel source of the spokes must be unique.

VTI guidelines

- When you configure the DVTIs of two hubs, ensure that they have the same IPsec tunnel mode (IPv4 or IPv6).
- Use distinct DVTIs for SD-WAN and route-based VPNs to avoid IPsec profile conflicts and errors.
- Configure the spokes' static VTIs in an ECMP zone to load balance the application traffic. If you do not configure the ECMP zone, the remaining paths act as backup paths when the primary path goes down. Note that you must configure the spokes' static VTIs and not the physical interfaces in the ECMP zone. This configuration is not part of the SD-WAN wizard.
- Ensure that you shutdown a DVTI virtual template before you change the IP address of a DVTI's numbered interface.

Limitations for using SD-WAN wizard

The SD-WAN wizard has these limitations and restrictions:

- You can configure a maximum of two hubs in an SD-WAN topology using the wizard.
- For each spoke, you can use only one WAN interface per topology. However, for dual-ISP setups, you can configure a second SD-WAN topology with the second WAN interface. For more information, see Sample configurations for dual-ISP deployment using SD-WAN wizard, on page 1797.
- SD-WAN wizard does not support the following:
  - IKEv1
  - Cluster devices are not supported on the hub and spoke because VTI is not supported on cluster devices.
  - Extranet hubs and spokes such as ASA, Cisco IOS, Cisco Viptela, Umbrella, Meraki, or vendor devices.

License requirements to use the SD-WAN wizard

Ensure that export-controlled features are enabled in your Smart License to configure an SD-WAN topology in Cloud-Delivered Firewall Management Center.

To verify if export-controlled functionality is enabled for your Smart License account, choose **System** (⚙) > **Licenses** > **Smart Licenses**.

Prerequisites for using SD-WAN wizard

Ensure that you review these prerequisites before you configure your SD-WAN topology using the SD-WAN wizard.

Device prerequisites

- Hub devices must be Version 7.6.0 and later.
- Spoke devices must be Version 7.3.0 and later.

General prerequisites

- You must be an Admin user.
- The Firewall Threat Defense devices must have an internet-routable public IP address. The IP address can be static or dynamic.
- Assign appropriate logical names and IP addresses to the interfaces of the Firewall Threat Defense devices. For example, use *inside* for the interface connected to the LAN, and *outside* for the interface connected to the internet or WAN.
- If you are using certificate-based authentication, you must enroll the certificates in the hub and spokes.
- Configure routing, NAT, and AC policies to ensure underlay connectivity between the devices.

Configure an SD-WAN topology using the SD-WAN wizard

The SD-WAN wizard allows you to easily configure VPN tunnels between your centralized headquarters and remote branch sites.

Before you begin

Ensure that you review Prerequisites for using SD-WAN wizard, on page 1792 and Guidelines for using SD-WAN wizard, on page 1790.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site** , and click **Add**.

Step 2 Enter a name for the SD-WAN VPN topology in the **Topology Name** field.

Step 3 Click the **SD-WAN Topology** radio button and click **Create**.

Step 4 Configure a hub:

a) Click **Add Hub**.

b) From the **Device** drop-down list, choose a hub.

c) Click **+** next to the **Dynamic Virtual Tunnel Interface (DVTI)** drop-down list to add a dynamic VTI for the hub.

The **Add Virtual Tunnel Interface** dialog box is prepopulated with default configurations. However, you must configure the **Tunnel Source**, and the **Borrow IP Address**. For more information, refer to Add a dynamic virtual tunnel interface for a hub, on page 1797.

d) In the **Hub Gateway IP Address** field, enter the public IP address of the hub's VPN interface or the tunnel source of the dynamic VTI to which the spokes connect.

This IP address is auto populated if the interface has a static IP address. If hub is behind a NAT device, you must manually configure the post-NAT IP address.

e) From the **Spoke Tunnel IP Address Pool** drop-down list, choose an IP address pool or click **+** to create an address pool.

When you add spokes, the wizard auto generates spoke tunnel interfaces, and assigns IP addresses to these spoke interfaces from this IP address pool.

f) Click **Add** to save the hub configuration.

g) (Optional) To add a secondary hub, repeat Step 4a to Step 4f.

h) Check the **ECMP** check box to enable Equal Cost Multi-Path (ECMP) on the dynamic VTIs of hub devices with Version 10.0 or later.

All virtual access interfaces on the hub connecting to the same spoke are grouped into an ECMP zone. This feature load balances traffic through multiple paths to a spoke.

Step 5 Configure spokes:

Click **Add Spoke** to add a single spoke device, or click **Add Spokes (Bulk Addition)** to add multiple spokes to your topology.

- Click **Add Spoke**. In the **Add Spoke** dialog box, configure the following parameters:

  a. From the **Device** drop-down list, choose a spoke.

  b. From the **VPN Interface** drop-down list, choose a WAN-facing or internet-facing physical interface to establish a VPN connection with the hub.

  c. Check the **Local Tunnel (IKE) Identity** check box to enable a unique and configurable identity for the VPN tunnel from this device to the remote peer. By default, this option is enabled.

  d. Choose one of the following options from the **Identity Type** drop-down list:

  - **Key ID**—(Default value) This value is auto populated as < sd-wan topologyname >_< device_IP_address >, for example, sdwantopo1_192.168.0.200. You can also specify a key ID of your choice.
  - **Email ID**—Specify an email ID up to 127 characters.
  - **IP Address**—IP address of the spoke's VPN interface.
  - **Auto**—IP address of the spoke's VPN interface for pre-shared key authentication or the certificate Distinguished Name (DN) for certificate-based authentication.
  - **Hostname**—Fully qualified hostname of the spoke.

  e. Click **Save** to save the spoke configuration.

- Click **Add Spokes (Bulk Addition)**. In the **Add Bulk Spokes** dialog box, configure the following parameters:

a. Choose one or more devices from the **Available Devices** list and click **Add** to move the devices to **Selected Devices** .

b. Use one of the following methods to select the VPN interfaces of the spokes:

- Click the **Interface Name Pattern** radio button and specify a string to match the logical name of the internet or WAN interface of the spokes, for example, outside*, wan*.

  If the spoke has multiple interfaces with the same pattern, the first interface that matches the pattern is selected for the topology.

- Click the **Security Zone** radio button and choose a security zone with the VPN interfaces of the spokes from the drop-down list, or click + to create a security zone.

c. Click **Next**.

The wizard validates if the spokes have interfaces with the specified pattern. Only the validated devices are added to the topology.

d. Click **Add**.

e. Click **Next**.

For each spoke, the wizard automatically selects the hub's DVTI as the tunnel source IP address.

Note
If the hub’s tunnel source IP address is an IPv6 address, the wizard automatically selects the first IPv6 address of the spokes' selected interface. To edit the IPv6 address of a spoke's tunnel source, click the edit icon next to a spoke, choose an IPv6 address from the **IP Address** drop-down list, and click **Save**.

Step 6 Configure authentication settings for the devices in the SD-WAN topology:

a) **Authentication Type**—For device authentication, you can use a manual pre-shared key, an auto-generated pre-shared key, or a certificate.

- **Pre-shared Manual Key**—Specify the pre-shared key for the VPN connection.
- **Pre-shared Automatic Key**—(Default value) The wizard automatically defines the pre-shared key for the VPN connection. Specify the key length in the **Pre-shared Key Length** field. The range is 1 to 127.
- **Certificate**—When you use certificates as the authentication method, the peers obtain digital certificates from a CA server in your PKI infrastructure, and use them to authenticate each other.

b) Choose one or more algorithms from the **Transform Sets** drop-down list.

c) Choose one or more algorithms from the **IKEv2 Policies** drop-down list.

Step 7 Configure the SD-WAN settings:

This step involves the auto generation of spoke tunnel interfaces, and BGP configuration of the overlay network.

a) From the **Spoke Tunnel Interface Security Zone** drop-down list, choose a security zone or click + to create a security zone to which the wizard automatically adds the spokes' auto-generated Static Virtual Tunnel Interfaces (SVTIs).

b) Check the **Enable BGP on the VPN Overlay Topology** check box to automate BGP configurations such as neighbor configurations between the overlay tunnel interfaces and basic route redistribution from the directly connected LAN interfaces of the hubs and spokes.

c) In the **Autonomous System Number** field, enter an Autonomous System (AS) number.

AS number is a unique number for a network with a single routing policy. BGP uses AS numbers to identify networks. The spoke's BGP neighbor configuration is generated based on the corresponding hub's AS number. Range is from 0 to 65536.

- If all the hubs and spokes are in the same region, by default, **64512** is the AS number.
- If the primary and secondary hubs are in different regions, the primary hub and the spokes are configured with **64512** as the AS number, and the secondary hub is configured with a different AS number.

d) In the **Community Tag for Local Routes** field, enter the BGP community attribute to tag connected and redistributed local routes. This attribute enables easy route filtering.

e) Check the **Redistribute Connected Interfaces** check box and choose an interface group from the drop-down list or click **+** to create an interface group with connected inside or LAN interfaces for BGP route redistribution in the overlay topology.

f) Check the **Enable Multiple Paths for BGP** check box to allow multiple BGP routes to be used at the same time to reach the same destination. This option enables BGP to load-balance traffic across multiple links.

Note that when you enable this option, BGP multipath is enabled only for spokes.

g) (Optional) Check the **Secondary Hub is in Different Autonomous System** check box. This check box appears only if you have a secondary hub in this topology.

h) In the **Autonomous System Number** field, enter the AS number for the secondary hub.

i) In the **Community Tag for Learned Routes** field, enter the BGP community attribute to tag routes learned from other SD-WAN peers over the VPN tunnel. This attribute is required only for eBGP configuration when the secondary hub has a different AS number. This field appears only if you have configured two hubs in the SD-WAN topology.

Step 8 Configure **Advanced Settings**.

a) From the **Identity Sent to Peers** drop-down list, choose the identity that the peers will use to identify themselves during IKE negotiations.

- **autoOrDN**—(Default value) Determines IKE negotiation by connection type: IP address for preshared key, or Cert DN for certificate authentication.
- **IP Address**—Uses the IP addresses of the hosts exchanging ISAKMP identity information.
- **Hostname**—Uses the fully qualified domain name of the hosts exchanging ISAKMP identity information. This name comprises the hostname and the domain name.

b) Configure **TrustSec (SGT) Settings**.

Enable SGT propagation over Virtual Tunnel Interfaces—Cisco TrustSec uses Security Group Tags (SGTs) to control access and enforce traffic on a network. This option enables SGT propagation over SVTIs and DVTIs of the VPN topology. To enable SGT propagation on a specific SVTI or DVTI, configure it in individual devices. Note that the Firewall Threat Defense device must be Version 10.0.0 and later.

c) Configure **Bidirectional Forwarding Detection (BFD) Settings**.

Enable Bidirectional Forwarding Detection Routing—BFD is a protocol for detecting forwarding path failures. This option enables BFD routing protocol on the SVTIs and DVTIs. When BFD detects a path failure, traffic is rerouted over the newly identified path. Note that the Firewall Threat Defense device must be Version 10.0.0 and later.

1. From the **Interval** drop-down list, choose the unit of interval, microseconds or milliseconds, at which BFD control packets are sent.
2. In the **Multiplier** field, enter the number of consecutive missed BFD control packets allowed before declaring that the peer is unavailable. The range is 3 to 50 packets.
3. In the **Minimum Transmit and Receive Interval** field, enter the minimum transmit and receive interval of the BFD control packets. The range is 50 to 999 milliseconds or 50000 to 999000 microseconds.

Step 9 Click **Finish** to save and validate the SD-WAN topology.

You can view the topology in the **Site-to-Site VPN Summary** page (**Devices** > **VPN** > **Site to Site**). After you deploy the configurations to all the devices, you can see the status of all the tunnels in this page.

What to do next

- View the auto-generated spoke SVTIs and their IP addresses—Click the edit icon next to the spoke configuration and click **View Generated Tunnel Interfaces** .
- We recommend that you enable ECMP on the spoke SVTIs. Choose **Devices** > **Device Management**, and click **Routing** > **ECMP**.
- Deploy the configurations on the hub and spokes. Choose **Deploy** . Select devices and click **Deploy** .
- Verify the SD-WAN topology tunnel statuses. For more information, see Verify tunnel statuses of an SD-WAN topology, on page 1802.
- Configure ACLs for the spokes' tunnel interface security zones. Choose **Policies** > **Access Control heading** > **Access Control**.
- We recommend that you enable BGP multipath on hubs. To enable BGP multipath on hubs:
  1. Choose ECMP **Devices** > **Device Management** , and click **Routing**.
  2. Under **General Settings**, click **BGP**.
  3. Check the **Enable BGP** check box to enable BGP.
  4. In the **AS Number** field, enter the AS number that you configured in the SD-WAN topology.
  5. Click **Save**.
  6. In the left pane, choose **BGP > IPv4 or IPv6**, and click the **General** tab.
  7. In the **Forward Packets over Multiple Paths** section, click the edit icon.
  8. Configure values for **Number of Paths** and **IBGP number of paths**. We recommend that you configure these values as 8.

- For more information about configuration examples using SD-WAN wizard, refer to Sample configurations for dual-ISP deployment using SD-WAN wizard, on page 1797.
- Configure a PBR policy on each spoke for application-aware routing based on the application performance metrics of the WAN interfaces. For more information, refer to Route Application Traffic from the Branch to the Internet Using Direct Internet Access (DIA).

Add a dynamic virtual tunnel interface for a hub

In the SD-WAN wizard, you must configure a DVTI for each hub. DVTI uses a virtual template to dynamically generate a unique virtual access interface for each VPN session.

Before you begin

In the SD-WAN wizard, click **Add Hub**, and choose a hub from the **Device** drop-down list.

Procedure

Step 1 Click + next to the **Dynamic Virtual Tunnel Interface (DVTI)** drop-down list to add a DVTI for the hub.

The **Add Virtual Tunnel Interface** dialog box appears with the following prepopulated default configurations.

a. **Tunnel Type**: Dynamic.

b. **Name**: <*tunnel_source_interface_logical name*>_dynamic_vti_<*tunnel_ID*>. For example, outside_dynamic_vti_1.

c. **Enabled** check box: Checked by default.

d. **Template ID**: Unique ID for the DVTI.

e. **Tunnel Source**: Physical interface that is the source of the DVTI and is auto populated by default.

f. **IPsec Tunnel Mode**: IPv4, by default.

Step 2 Choose a security zone for the dynamic VTI from the **Security Zone** drop-down list.

Step 3 Choose a physical or loopback interface from the **Borrow IP** drop-down list, the dynamic VTI interface inherits this IP address.

Ensure that you use an IP address different from the tunnel source IP address. We recommend that you use a loopback IP address.

Step 4 Click **OK** to save the dynamic VTI.

Sample configurations for dual-ISP deployment using SD-WAN wizard

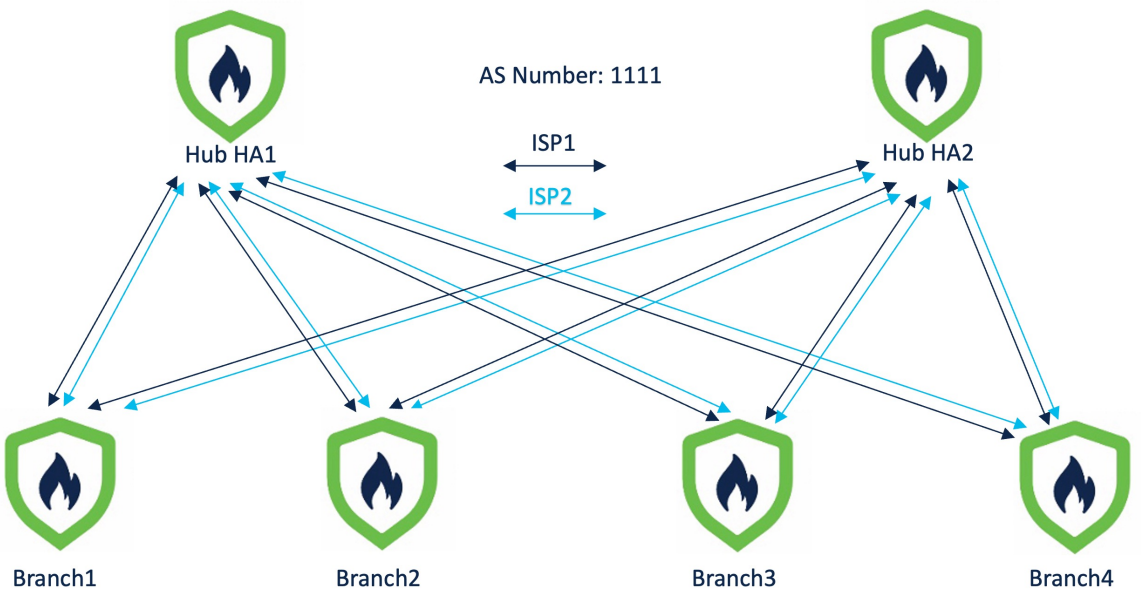
Here are some sample configurations that demonstrate how to implement dual-ISP deployment using the SD-WAN wizard for redundancy and load balancing across multiple internet service providers.

Dual ISP deployment: Two hubs and four spokes in the same region

In this dual ISP topology, the hubs and the spokes are in a single region, with AS number AS 1111. The hubs and spokes use Internal Border Gateway Protocol (iBGP) as the routing protocol to exchange routing information.

- Hub HA1 and Hub HA2 are hub Threat Defense devices at the headquarters.
- Branch1, Branch2, Branch3, and Branch4 are spoke Threat Defense devices at the branches.
- ISP1 is the VPN interface of each spoke to ISP1.
- ISP2 is the VPN interface of each spoke to ISP2.

Figure 372: Dual ISP topology with two hubs and four spokes in the same region



To configure this topology, you must create two SD-WAN topologies using the SD-WAN wizard.

SD-WAN topology 1

Parameter	Value
Primary Hub	Hub HA1
Secondary Hub	Hub HA2
Spokes	Branch1, Branch2, Branch3, Branch4
AS Number	1111
VPN Interface (Spoke Tunnel Source)	ISP1
Number of Tunnels	8

The total number of tunnels in SD-WAN Topology 1 is 8.

SD-WAN topology 2

Parameter	Value
Primary Hub	Hub HA1
Secondary Hub	Hub HA2
Spokes	Branch1, Branch2, Branch3, Branch4
AS Number	1111
VPN Interface (Spoke Tunnel Source)	ISP2
Number of Tunnels	8

The total number of tunnels in SD-WAN Topology 2 is 8.

The total number of VPN tunnels for this dual ISP deployment is 16.

Note If the hubs are in different geographic locations and have different protected networks behind them, to ensure direct communication between these networks, configure a point-to-point route-based VPN topology between the two hubs using the route-based VPN wizard.

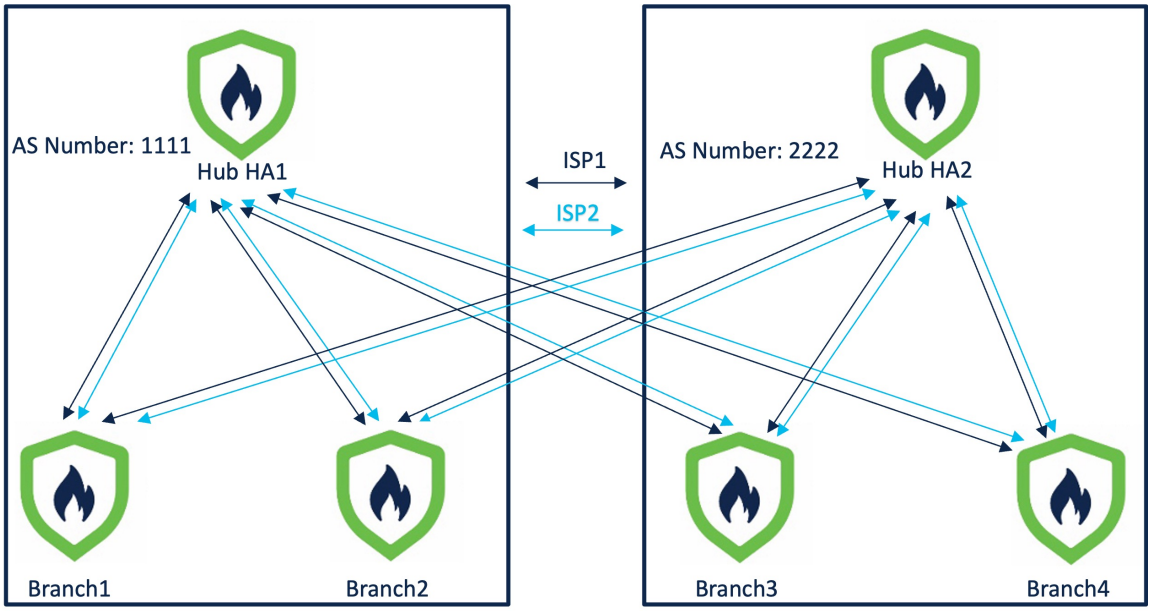
Dual ISP deployment: Two hubs and four spokes in different regions

In this dual ISP topology, the hubs are in different regions, and have two directly connected spokes each. The hubs and their directly connected spokes use Internal Border Gateway Protocol (iBGP) as the routing protocol, and the hubs use External Border Gateway Protocol (eBGP) to exchange routing information.

Topology components

- Hub HA1 and Hub HA2 are hub Threat Defense devices at the headquarters.
- Branch1, Branch2, Branch3, and Branch4 are spoke Threat Defense devices at the branches.
- HQ1, Branch1, and Branch2 are in a single region with AS number AS 1111.
- HQ2, Branch3, and Branch4 are in a single region with AS number AS 2222.
- ISP1 is the VPN interface of each spoke to ISP1.
- ISP2 is the VPN interface of each spoke to ISP2.

Figure 373: Dual ISP topology with two hubs and four spokes in different regions



To configure this topology, you must create the four SD-WAN topologies using the SD-WAN wizard:

SD-WAN topology 1

Parameter	Value
Primary Hub	Hub HA1
Secondary Hub	Hub HA2
Spokes	Branch1, Branch2
AS Number	1111
Secondary AS Number	2222
VPN Interface (Spoke Tunnel Source)	ISP1

The number of tunnels in SD-WAN Topology 1 is 4.

SD-WAN topology 2

Parameter	Value
Primary Hub	Hub HA1
Secondary Hub	Hub HA2
Spokes	Branch1, Branch2
AS Number	1111

Parameter	Value
Secondary AS Number	2222
VPN Interface (Spoke Tunnel Source)	ISP2

The number of tunnels in SD-WAN Topology 2 is 4.

SD-WAN topology 3

Parameter	Value
Primary Hub	Hub HA2
Secondary Hub	Hub HA1
Spokes	Branch3, Branch4
AS Number	2222
Secondary AS Number	1111
VPN Interface (Spoke Tunnel Source)	ISP1

The number of tunnels in SD-WAN Topology 3 is 4.

SD-WAN topology 4

Parameter	Value
Primary Hub	Hub HA2
Secondary Hub	Hub HA1
Spokes	Branch3, Branch4
AS Number	2222
Secondary AS Number	1111
VPN Interface (Spoke Tunnel Source)	ISP2

The number of tunnels in SD-WAN Topology 4 is 4.

The total number of VPN tunnels for this dual ISP deployment is 16.



Note If the hubs are in different geographic locations and have different protected networks behind them, to ensure direct communication between these networks, configure a point-to-point route-based VPN topology between the two hubs using the route-based VPN wizard.

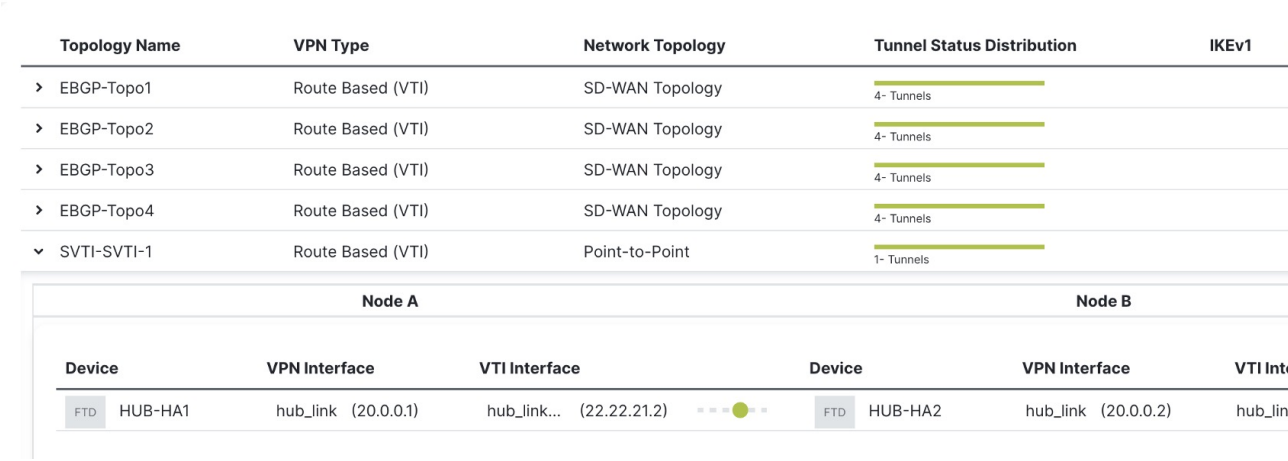
Verify tunnel statuses of an SD-WAN topology

You can verify the tunnel statuses of SD-WAN topologies, including checking VPN tunnel status in the Site-to-Site VPN Summary page and Dashboard, viewing virtual tunnel interfaces on devices, and verifying BGP routing configuration on hubs and branches.

Verify tunnel statuses on the Site-to-Site VPN summary page

To verify if the VPN tunnels of the SD-WAN topologies are up, choose **Devices** > **VPN** > **Site to Site**.

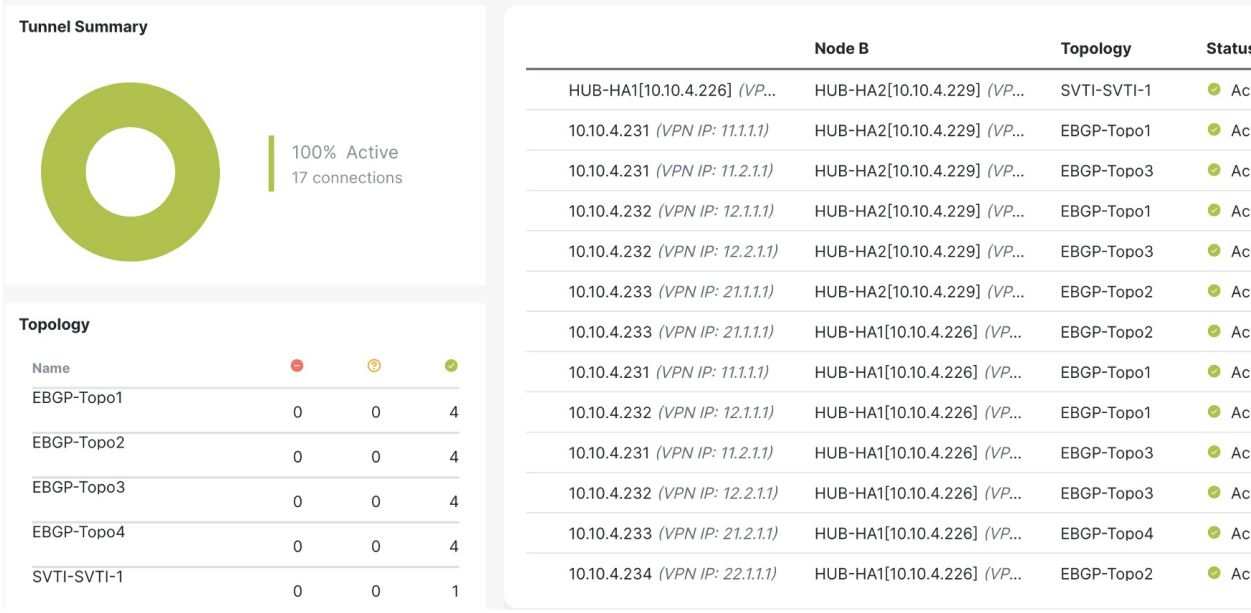
This example shows the five SD-WAN topologies with two hubs and four spokes in different regions that are connected to dual ISPs:



Verify tunnel statuses on the Site-to-Site VPN dashboard

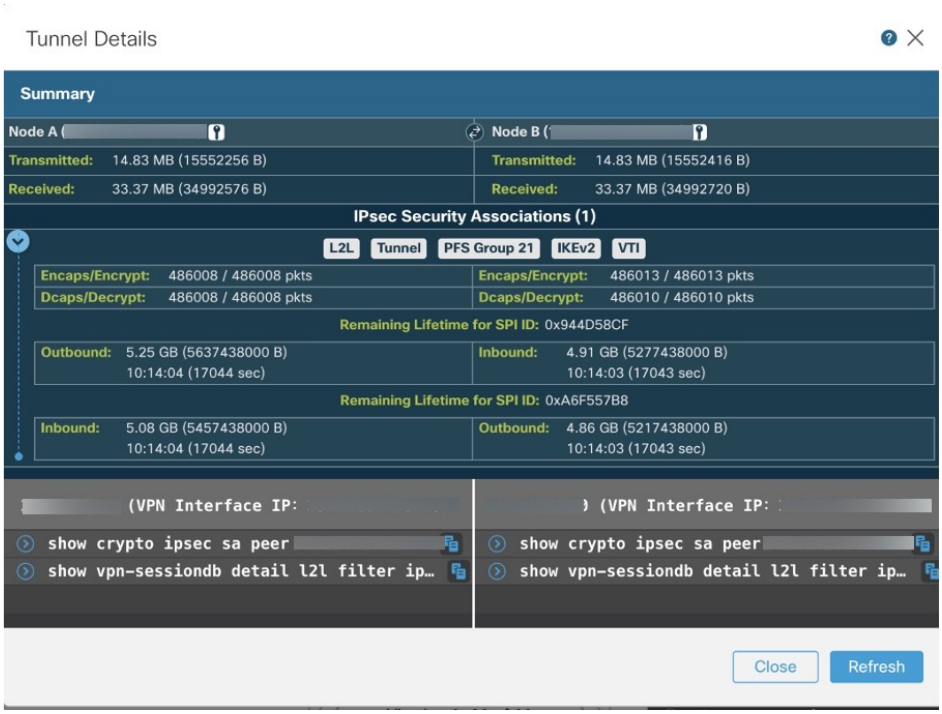
To view details of the SD-WAN VPN tunnels, choose **Overview** > **Dashboards** > **Site to Site VPN**.

This example shows the VPN tunnels for an SD-WAN topology with two hubs and four spokes in different regions that are connected to dual ISPs:



To view more details of each VPN tunnel:

1. Hover over a tunnel.
2. Click the **View Full Information** ( ) icon. A pane with tunnel details and more actions is displayed.
3. Click the **CLI Details** tab in the side pane to view the show commands and details of the IPsec security associations.



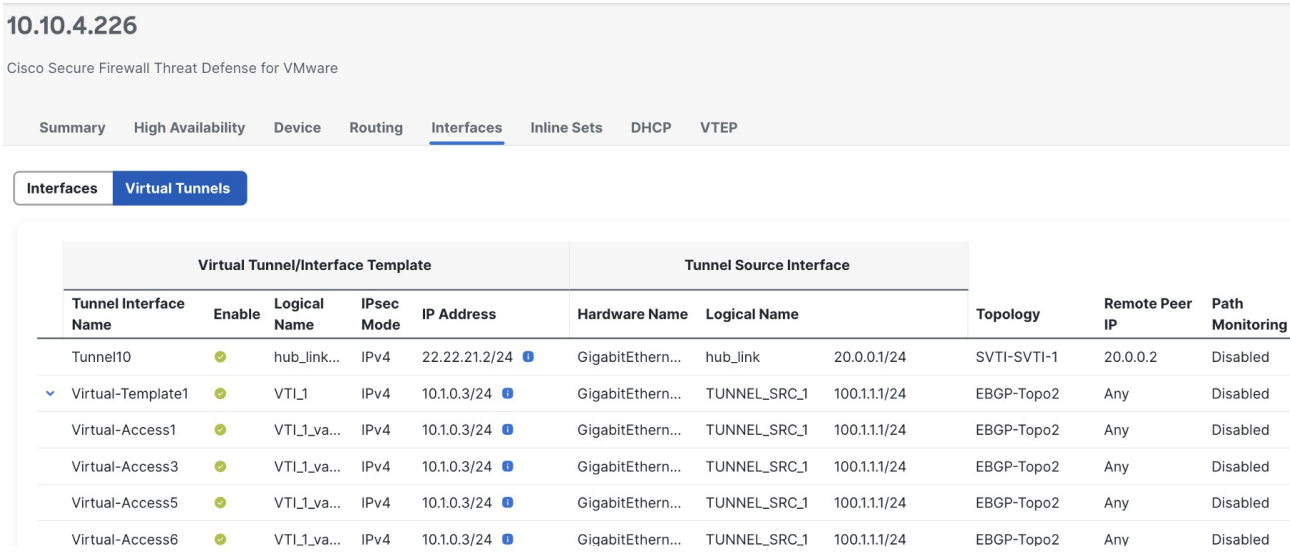
View virtual tunnel interfaces of the devices

To view the dynamic VTIs of hubs and static VTIs of spokes:

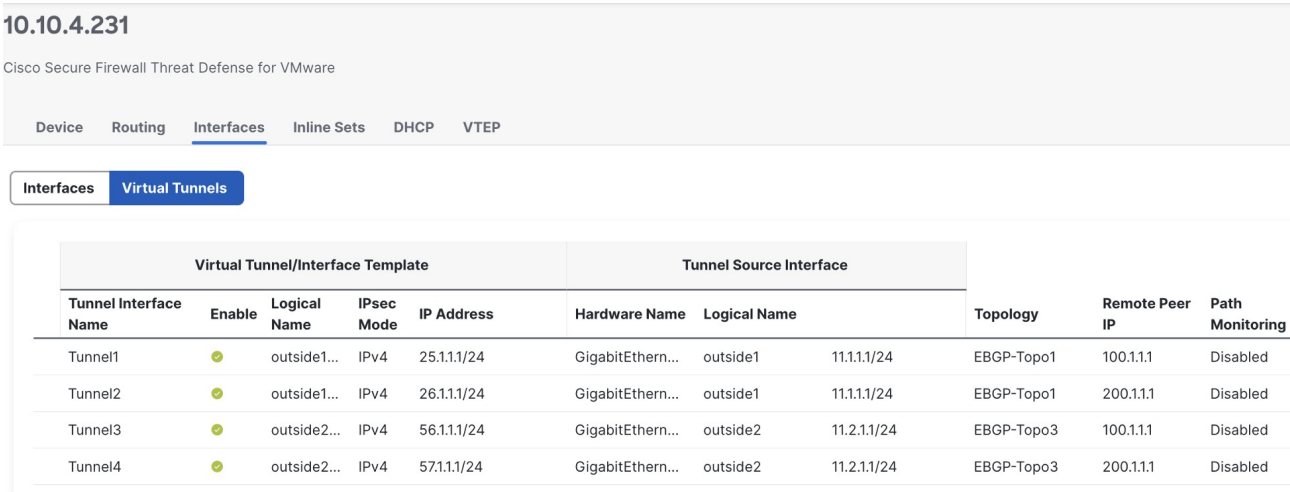
1. Choose **Devices** > **Device Management**.
2. Click the edit icon for a hub or a spoke device.
3. Click the **Interface** tab.
4. Click the **Virtual Tunnels** tab.

   For each VTI, you can view details such as name, IP address, IPsec mode, tunnel source interface details, topology, and remote peer IP.

This image shows an example of the virtual access interfaces created dynamically by a hub's DVTI:



This image shows an example of the static tunnel virtual interfaces (SVTIs) created on a spoke by the SD-WAN wizard:



The SD-WAN wizard assigns IP addresses to these tunnel interfaces from the IP address pool of the hub.

Verify routing in the hub and branch devices

To verify the BGP configuration of the hubs and spokes of the SD-WAN topologies:

1. Choose **Devices** > **Device Management**.
2. Click the edit icon for a hub or a spoke device.
3. Click the **Device** tab.
4. Click **CLI** in the **General** card. The **CLI Troubleshoot** window appears.
5. Enter one of these commands in the **Command** field and click **Execute**:
   - **show route**
   - **show BGP summary**

Monitor SD-WAN topologies

This section provides information about the different ways to monitor your SD-WAN topologies.

SD-WAN summary dashboard

The SD-WAN Summary dashboard (**Overview** > **Dashboards** > **SD-WAN SummaryAnalysis** > **SD-WAN Summary**) provides a snapshot of your WAN devices and their interfaces. This dashboard helps you to:

- Identify issues with the underlay and overlay (VPN) topologies.
- Troubleshoot VPN issues using the existing **Health Monitoring**, **Device Management**, and **Site-to-Site Monitoring** pages.
- Monitor application performance metrics of WAN interfaces. The threat defense steers application traffic based on these metrics.

For clusters, this dashboard displays application performance metrics of only the control node and not the data nodes.

WAN device criteria

A WAN device must meet one of these criteria:

- The device must be a VPN peer.
- The device must have WAN interface.

WAN interface criteria

A WAN interface must meet one of these criteria:

- The interface has IP address-based path monitoring enabled on it.

- The interface has a Policy Based Routing (PBR) policy with at least one application configured to monitor it.

For more information about PBR policy and path monitoring, see Policy-Based Routing, on page 1541.

Uplink decisions

Click **Uplink Decisions** to view the **VPN Troubleshooting** page. You can view syslogs with ID: 880001. These syslogs show the threat defense interfaces through which it steers traffic based on the configured PBR policy.

To view the above syslogs and to view the data on this dashboard, ensure that you review Prerequisites for using SD-WAN summary dashboard, on page 1806.

Prerequisites for using SD-WAN summary dashboard

Ensure that you review these prerequisites before using the SD-WAN summary dashboard.

General prerequisites

- You must be an Admin, Security Analyst, or Maintenance user to view this dashboard. See Secure Firewall Management Center and Cloud-delivered Firewall Management Center User Role Mapping for more information.
- Threat Defense devices must be Version 7.2 or later.
- Enable IP-based path monitoring and HTTP-based application monitoring on the WAN interfaces.

  To do this:

  1. Choose **Devices** > **Device Management**.
  2. Click the edit icon adjacent to the device that you want to edit.
  3. Click the edit icon adjacent to the interface that you want to edit.
  4. Click the **Path Monitoring** tab.
  5. Check the **Enable IP based Monitoring** check box.
  6. Check the **Enable HTTP based Application Monitoring** check box.
  7. Click **OK**.

- Configure a PBR policy with at least one application configured to monitor it:

  1. Choose **Devices** > **Device Management**.
  2. Click the edit icon adjacent to the device that you want to edit.
  3. Click **Routing**.
  4. In the left pane, click **Policy Based Routing**.
  5. Click **Add**.
  6. From the **Ingress Interface** drop-down list, choose an interface.
  7. Click **Add** to configure a forwarding action.

8. Configure the parameters.
9. Click **Save**.

- To view syslogs when you click **Uplink Decisions**, you must:
  - Choose **Devices** > **Platform Settings** and create or edit a threat defense policy.
  - In the left pane, click **Syslog**.
  - Click the **Logging Setup** tab.
  - Check the **Enable Logging** check box to turn on the data plane system logging for the threat defense device.
  - Click the **All Logs** radio button to enable logging of all the troubleshooting syslog messages.

    or

    Click the **VPN Logs** radio button to enable logging of only the VPN troubleshooting messages.
  - Click **Save**.

Monitor WAN devices and interfaces using the SD-WAN summary dashboard

The SD-WAN Summary dashboard has these widgets :

- WAN connectivity, on page 1807
- VPN topology, on page 1808
- interface throughput, on page 1808
- Device inventory, on page 1808
- WAN device health, on page 1808

WAN connectivity

This widget provides a summary of the WAN interfaces statuses. It shows the number of WAN interfaces that are in the **Online**, **Offline** or **No Data** states. Note that you cannot monitor subinterfaces using this widget.

Click **View All Interfaces** to view more details about the interfaces in the health monitor page.

If a WAN interface is in the **Offline** or **No Data** state, you can troubleshoot it from the health monitor page:

1. In the **Monitoring** pane, expand **Devices**.
2. Click the corresponding WAN device to view the device-specific health details.
3. Click the **Interface** tab to view the interface status and aggregate traffic statistics for a specific time.

   Alternatively, you can click **View System & Troubleshoot Details**. The health monitor page is displayed with all the necessary details.

VPN topology

This widget provides a summary of the site-to-site VPN tunnel statuses. It shows the number of **Active**, **Inactive**, and **No Active Data** VPN tunnels.

Click **View All Connections** to view the VPN tunnel details in the **Site-to-site VPN Monitoring** dashboard.

If the tunnels are in the **Inactive** or **No Active Data** state, you can troubleshoot using the **Site-to-site VPN Monitoring** dashboard. In the **Tunnel Status** widget, hover your cursor over a topology, click **View** ( ) and do one of these actions:

- Click the **CLI Details** tab to view the details of the VPN tunnels.
- Click the **Packet Tracer** tab to use the packet tracer tool for the topology.

interface throughput

This widget monitors the average throughput of the WAN interfaces during the chosen time period.

The interface throughput is classified into four bands. These details aid in cost planning and resourcing. You can choose a time range for the widget data from the **Show Last** drop-down list. The range is from 15 minutes to two weeks.

Click **View Health Monitoring** to view more details about the interface in the health monitor page.

Device inventory

This widget lists all the managed WAN devices and groups them according to the model.

Click **View Device Management** to view more details about the device in the **Device Management** page.

WAN device health

This widget displays the device count according to the health of the WAN devices. You can view the number of devices with errors, warnings, or those that are in **Disabled** state.

Click **View Health Monitoring** to view the alarms, and quickly identify, isolate, and resolve issues.

If the health of a device is affected, you can troubleshoot it from the health monitor page.

1. In the **Monitoring** pane, expand **Devices**.
2. Click the corresponding WAN device to view the device-specific health details.
3. Click **View System & Troubleshoot Details**. The health monitor page is displayed with all the necessary details.

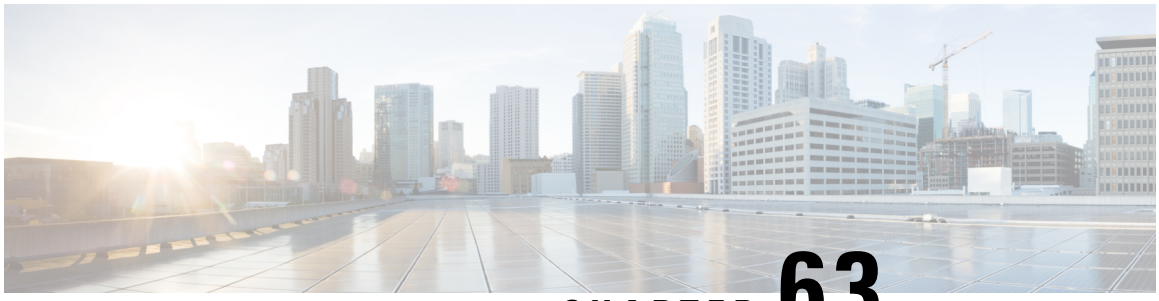
A device can be in **Disabled** state for multiple reasons, including these:

- Management interface is disabled.
- Device is powered off.
- Device is being upgraded.

Monitor application performance metrics of WAN interfaces using SD-WAN summary dashboard

You can monitor application performance metrics of WAN interfaces using the **Application Monitoring** tab of the SD-WAN Summary dashboard. These metrics include Jitter, Round Trip Time (RTT), Mean Opinion Score (MOS), and Packet Loss.

By default, the metrics data is refreshed every 5 minutes. You can change the refresh time; the range is from 5 to 30 minutes. You can view the metrics in tabular and graphical formats. For each WAN interface, the latest metric value appears in the table. For graphical data, you can choose a time interval of up to 24 hours to view the metrics data for the corresponding WAN interfaces.



CHAPTER 63

Remote Access VPN

Remote Access virtual private network (VPN) allows individual users to connect to your network from a remote location using a computer or other supported devices connected to the Internet. This allows mobile workers to connect from their home networks or a public Wi-Fi network, for example.

The following topics explain how to configure remote access VPN for your network.

- Remote Access VPN Overview, on page 1811
- License Requirements for Remote Access VPN, on page 1817
- Requirements and Prerequisites for Remote Access VPN, on page 1817
- Guidelines and Limitations for Remote Access VPNs, on page 1818
- Configuring a New Remote Access VPN Connection, on page 1821
- Create a Copy of an Existing Remote Access VPN Policy, on page 1831
- Set Target Devices for a Remote Access VPN Policy, on page 1831
- Associate Local Realm with Remote Access VPN Policy, on page 1832
- Additional Remote Access VPN Configurations, on page 1832
- Customizing Remote Access VPN AAA Settings, on page 1890
- Advanced Secure Client Configurations, on page 1910
- Remote Access VPN Examples, on page 1919
- Monitor Remote Access VPNs, on page 1924

Remote Access VPN Overview

Secure Firewall Threat Defense provides secure gateway capabilities that support remote access SSL and IPsec-IKEv2 VPNs. The full tunnel client, Secure Client, provides secure SSL and IPsec-IKEv2 connections to the security gateway for remote users. When the client negotiates an SSL VPN connection with Firewall Threat Defense, it connects using Transport Layer Security (TLS) or Datagram Transport Layer Security (DTLS)

Secure Client is the only client supported on endpoint devices for remote VPN connectivity to Firewall Threat Defense devices. The client gives remote users the benefits of an SSL or IPsec-IKEv2 VPN client without the need for network administrators to install and configure clients on remote computers. The Secure Client for Windows, Mac, and Linux is deployed from the secure gateway upon connectivity. The Secure Client apps for Apple iOS and Android devices are installed from the platform app store.

Use the Remote Access VPN policy wizard to set up SSL and IPsec-IKEv2 remote access VPNs with basic capabilities. Then, enhance the policy configuration as you want and deploy it to your Firewall Threat Defense secure gateway devices.

Remote Access VPN Features

The following table describes the features of Secure Firewall Threat Defense remote access VPN:

Table 103: Remote access VPN features

	Description
Secure Firewall Threat Defense remote access VPN features	• SSL and IPsec-IKEv2 remote access using the Secure Client. • Cloud-Delivered Firewall Management Center supports all combinations such as IPv6 over an IPv4 tunnel. • Configuration support on both Cloud-Delivered Firewall Management Center and Firewall Device Manager. Device-specific overrides. • Support for both Cloud-Delivered Firewall Management Center and Firewall Threat Defense HA environments. • Support for multiple interfaces and multiple AAA servers. • Rapid Threat Containment support using RADIUS CoA or RADIUS dynamic authorization. • Support for DTLS v1.2 protocol with Cisco Secure Client version 4.7 or higher. • Secure Client modules support for additional security services for remote access VPN connections. • VPN load balancing.

	Description
AAA features	• Server authentication using self-signed or CA-signed identity certificates. • AAA username and password-based remote authentication using RADIUS server or LDAP or AD. • RADIUS group and user authorization attributes, and RADIUS accounting. • Double authentication support using an additional AAA server for secondary authentication. • NGFW Access Control integration using VPN Identity. • LDAP or AD authorization attributes using Cloud-Delivered Firewall Management Center web interface. • Support for single sign-on using SAML 2.0. • Support for multiple identity provider trustpoints with Microsoft Azure that can have multiple applications for the same Entity ID, but a unique identity certificate. • Restrict remote access VPN connections based on their geolocations.
VPN tunneling features	• Address assignment. • Split tunneling. • Split DNS. • Client Firewall ACLs. • Session Timeouts for maximum connect and idle time.
Remote access VPN monitoring features	• New VPN Dashboard Widget showing VPN users by various characteristics such as duration and client application. • Remote access VPN events including authentication information such as username and OS platform. • Tunnel statistics available using the Firewall Threat Defense Unified CLI.

Secure Client Components

Secure Client Deployment

Your remote access VPN policy can include the Secure Client Image and the Secure Client Profile for distribution to connecting endpoints. Or, the client software can be distributed using other methods. See the *Deploy Cisco Secure Client* chapter in the Cisco Secure Client (including AnyConnect) Administrator Guide, Release 5.

Without a previously installed client, remote users enter the IP address in their browser of an interface configured to accept SSL or IPsec-IKEv2 VPN connections. Unless the security appliance is configured to redirect http:// requests to https://, remote users must enter the URL in the form https://*address*. After the user enters the URL, the browser connects to that interface and displays the login screen.

After a user logs in, if the secure gateway identifies the user as requiring the VPN client, it downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure connection, and either remains or uninstalls itself (depending on the security appliance configuration) when the connection stops. In the case of a previously installed client, after login, the Firewall Threat Defense security gateway examines the client version and upgrades it as necessary.

Secure Client Operation

When the client negotiates a connection with the security appliance, the client connects using Transport Layer Security (TLS), and optionally, Datagram Transport Layer Security (DTLS). DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.

When an IPsec-IKEv2 VPN client initiates a connection to the secure gateway, negotiation consists of authenticating the device through Internet Key Exchange (IKE), followed by user authentication using IKE Extended Authentication (Xauth). The group profile is pushed to the VPN client and an IPsec security association (SA) is created to complete the VPN.

Secure Client Profile and Editor

The Secure Client Profile is a group of configuration parameters, stored in an XML file that the VPN client uses to configure its operation and appearance. These parameters (XML tags) include the names and addresses of host computers and settings to enable more client features.

You can configure a profile using the Secure Client Profile Editor. This editor is a convenient GUI-based configuration tool that is available as part of the Secure Client software package. It is an independent program that you run outside of the Cloud-Delivered Firewall Management Center.

Remote Access VPN Authentication

Remote Access VPN Server Authentication

Secure Firewall Threat Defense secure gateways always use certificates to identify and authenticate themselves to the VPN client endpoint.

While you use the Remote Access VPN Policy Wizard, you can enroll the selected certificate on the targeted Firewall Threat Defense device. In the wizard, under **Access & Certificate** phase, select “Enroll the selected certificate object on the target devices” option. The certificate enrollment gets automatically initiated on the specified devices. As you complete the remote access VPN policy configuration, you can view the status of

the enrolled certificate under the device certificate homepage. The status provides a clear standing as to whether the certificate enrollment was successful or not. Your remote access VPN policy configuration is now fully completed and ready for deployment.

Obtaining a certificate for the secure gateway, also known as PKI enrollment, is explained in Certificates, on page 3079. This chapter contains a full description of configuring, enrolling, and maintaining gateway certificates.

Remote Access VPN Client AAA

For both SSL and IPsec-IKEv2, remote user authentication is done using usernames and passwords only, certificates only, or both.



Note If you are using client certificates in your deployment, they must be added to your client's platform independent of the Secure Firewall Threat Defense or Cloud-Delivered Firewall Management Center. Facilities such as SCEP or CA Services are not provided to populate your clients with certificates.

AAA servers enable managed devices acting as secure gateways to determine who a user is (authentication), what the user is permitted to do (authorization), and what the user did (accounting). Some examples of the AAA servers are RADIUS, LDAP/AD, TACACS+, and Kerberos. For Remote Access VPN on Firewall Threat Defense devices, AD, LDAP, and RADIUS AAA servers are supported for authentication.

Refer to the section Understanding Policy Enforcement of Permissions and Attributes to understand more about remote access VPN authorization.

Before you add or edit the remote access VPN policy, you must configure the Realm and RADIUS server groups you want to specify. For more information, see Create an LDAP realm or an Active Directory realm and realm directory, on page 2325 and Add a RADIUS server group, on page 2943.

Without DNS configured, the device cannot resolve AAA server names, named URLs, and CA Servers with FQDN or Hostnames, it can only resolve IP addresses.

The login information provided by a remote user is validated by an LDAP or AD realm or a RADIUS server group. These entities are integrated with the Secure Firewall Threat Defense secure gateway.



Note If users authenticate with remote access VPN using Active Directory as the authentication source, users must log in using their username; the format `domain\username` or `username@domain` fails. (Active Directory refers to this username as the *logon name* or sometimes as `sAMAccountName`.) For more information, see User Naming Attributes on MSDN.

If you use RADIUS to authenticate, users can log in with any of the preceding formats.

Once authenticated via a VPN connection, the remote user takes on a *VPN Identity*. This VPN Identity is used by *identity policies* on the Secure Firewall Threat Defense secure gateway to recognize and filter network traffic belonging to that remote user.

Identity policies are associated with access control policies, which determine who has access to network resources. It is in this way that the remote user blocked or allowed to access your network resources.

For more information, see the Identity policies, on page 2505 and Access Control Policies, on page 1997 sections.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1835

Understanding Policy Enforcement of Permissions and Attributes

The Secure Firewall Threat Defense device supports applying user authorization attributes (also called user entitlements or permissions) to VPN connections from an external authentication server and/or authorization AAA server (RADIUS) or from a group policy on the Firewall Threat Defense device. If the Firewall Threat Defense device receives attributes from the external AAA server that conflicts with those configured on the group policy, then attributes from the AAA server always take the precedence.

The Firewall Threat Defense device applies attributes in the following order:

1. **User attributes on the external AAA server**—The server returns these attributes after successful user authentication and/or authorization.
2. **Group policy configured on the Firepower Threat Defense device**—If a RADIUS server returns the value of the RADIUS Class attribute IETF-Class-25 (OU= group-policy) for the user, the Firewall Threat Defense device places the user in the group policy of the same name and enforces any attributes in the group policy that are not returned by the server.
3. **Group policy assigned by the Connection Profile (also known as Tunnel Group)**—The Connection Profile has the preliminary settings for the connection, and includes a default group policy applied to the user before authentication.



Note The Firewall Threat Defense device does not support inheriting system default attributes from the default group policy, *DfltGrpPolicy*. The attributes on the group policy assigned to the connection profile are used for the user session, if they are not overridden by user attributes or the group policy from the AAA server as indicated above.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1835

Understanding AAA Server Connectivity

LDAP, AD, and RADIUS AAA servers must be reachable from the Firewall Threat Defense device for your intended purposes: user-identity handling only, VPN authentication only, or both activities. AAA servers are used in remote access VPN for the following activities:

- **User-identity handling**— the servers must be reachable over the Management interface.

  On the Firewall Threat Defense, the Management interface has a separate routing process and configuration from data interfaces.

- **VPN authentication**—the servers must be reachable over a data interface or the Management interface.

  To use the Management interface, you must explicitly select Management as the source interface. Other management-only interfaces cannot be used.

To use the same AAA servers for both activities, we recommend specifying the Management interface as the source interface.

For more information about various interfaces, see Regular Firewall Interfaces, on page 1039.

After deployment, use the following CLI commands to monitor and troubleshoot AAA server connectivity from the Firewall Threat Defense device:

- **show aaa-server** to display AAA server statistics.
- **show network** and **show network-static-routes** to view the Management interface default route and static routes.
- **show route** to view data traffic routing table entries.
- **ping system** and **traceroute system** to verify the path to the AAA server through the Management interface.
- **ping interface** *ifname* and **traceroute** *destination* to verify the path to the AAA server through the data interfaces.
- **test aaa-server authentication** and **test aaa-server authorization** to test authentication and authorization on the AAA server.
- **clear aaa-server statistics** *groupname* or **clear aaa-server statistics protocol** *protocol* to clear AAA server statistics by group or protocol.
- **aaa-server** *groupname* **active host** *hostname* to activate a failed AAA server, or **aaa-server** *groupname* **fail host** *hostname* to fail a AAA server.
- **debug ldap level**, **debug aaa authentication**, **debug aaa authorization**, and **debug aaa accounting**.

License Requirements for Remote Access VPN

Threat Defense License

Firewall Threat Defense remote access VPN requires Strong Encryption and one of the following licenses for Secure Client:

- Secure Client Advantage
- Secure Client Premier
- Secure Client VPN Only

Requirements and Prerequisites for Remote Access VPN

Model support

Firewall Threat Defense

Supported domains

Any

User roles

Admin

Guidelines and Limitations for Remote Access VPNs

Remote Access VPN Policy Configuration

- You can add a new remote access VPN policy only by using the wizard. You must proceed through the entire wizard to create a new policy; the policy will not be saved if you cancel before completing the wizard.
- Two users must **not** edit a remote access VPN policy at the same time; however, the web interface does not prevent simultaneous editing. If this occurs, the last saved configuration persists.
- Moving a Secure Firewall Threat Defense device from one domain to another domain is not possible if remote access VPN policy is assigned to that device.
- Remote access VPN does not support SSL while using ECMP. We recommend that you use IPsec-IKEv2.
- Remote access VPN does not support clustering.
- Remote access VPN connectivity could fail if there is a misconfigured Firewall Threat Defense NAT rule.
- If you are using DHCP to provide IP addresses to the client, and the client cannot obtain an address, check the NAT rules. Any NAT rule that applies to the RA VPN network should include the route lookup option. Route lookup can help ensure the DHCP requests are sent to the DHCP server through an appropriate interface.
- Whenever IKE ports 500/4500 or SSL port 443 is in use or when there are some PAT translations that are active, the Secure Client IPSec-IKEv2 or SSL remote access VPN cannot be configured on the same port as it fails to start the service on those ports. These ports must not be used on the Firewall Threat Defense device before configuring remote access VPN policy.
- While configuring remote access VPNs using the wizard, you can create in-line certificate enrollment objects, but you cannot use them to install the identity certificate. Certificate enrollment objects are used for generating the identity certificate on the Firewall Threat Defense device being configured as the remote access VPN gateway. Install the identity certificate on the device before deploying the remote access VPN policy to the device.

  For more information about how to install the identity certificate based on the certificate enrollment object, see Use the object manager, on page 2932.
- The ECMP zone interfaces can be used in remote access VPN with IPsec enabled.
- The ECMP zone interfaces cannot be used in remote access VPN with SSL enabled. Deployment of remote access VPN (SSL enabled) configuration fails if all the remote access VPN interfaces that belong to security zones or interface groups also belong to one or more ECMP zones. However, if only some of the remote access VPN interfaces belonging to the security zones or interface groups also belongs to one or more ECMP zones, deployment of the remote access VPN configuration succeeds excluding those interfaces.
- After you change the remote access VPN policy configurations, re-deploy the changes to the Firewall Threat Defense devices. The time it takes to deploy configuration changes depends on multiple factors such as complexity of the policies and rules, type and volume of configurations you send to the device, and memory and device model. Before deploying remote access VPN policy changes, review the Best practices for deploying configuration changes, on page 200.

- Issuing commands such as **curl** against the RA VPN headend is not directly supported, and might not have desirable results. For example, the headend does not respond to HTTP HEAD requests.
- Firewall Threat Defense does not accept remote access VPN sessions if the third-party client sends a null user agent.
- You can configure browser proxy using FlexConfig.
- When **Bypass Access Control policy for decrypted traffic** is disabled on a Firewall Threat Defense and ISE issues a DACL, traffic is first matched against the ACP: if permitted, it proceeds to DACL checking; if denied, it bypasses DACL matching.

Concurrent VPN Sessions Capacity Planning (Firewall Threat Defense Virtual Models)

The maximum concurrent VPN sessions are governed by the installed Firewall Threat Defense Virtual smart-licensed entitlement tier, and enforced via a rate limiter. There is a maximum limit to the number of concurrent remote access VPN sessions allowed on a device based on the licensed device model. This limit is designed so that system performance does not degrade to unacceptable levels. Use these limits for capacity planning.

Device Model	Maximum Concurrent Remote Access VPN Sessions
Firewall Threat Defense Virtual5	50
Firewall Threat Defense Virtual10	250
Firewall Threat Defense Virtual20	250
Firewall Threat Defense Virtual30	250
Firewall Threat Defense Virtual50	750
Firewall Threat Defense Virtual100	10,000

Concurrent VPN Sessions Capacity Planning (Hardware Models)

The maximum concurrent VPN sessions are governed by platform-specific limits and have no dependency on the license. There is a maximum limit to the number of concurrent remote access VPN sessions allowed on a device based on the device model. This limit is designed so that system performance does not degrade to unacceptable levels. Use these limits for capacity planning.

Device Model	Maximum Concurrent Remote Access VPN Sessions
Firepower 1010	75
Firepower 1120	150
Firepower 1140	400
Firepower 2110	1500
Firepower 2120	3500
Firepower 2130	7500

Device Model	Maximum Concurrent Remote Access VPN Sessions
Firepower 2140	10,000
Secure Firewall 3110	3000
Secure Firewall 3120	6000
Secure Firewall 3130	15,000
Secure Firewall 3140	20,000
Firepower 4100, all models	10,000
Firepower 9300 appliance, all models	20,000
ISA 3000	25

For capacity of other hardware models, contact your sales representative.



Note The Firewall Threat Defense device denies the VPN connections once the maximum session limit per platform is reached. The connection is denied with a syslog message. Refer the syslog messages %ASA-4-113029 and %ASA-4-113038 in the syslog messaging guide. For more information, see Cisco Secure Firewall ASA Series Syslog Messages.

Controlling Cipher Usage for VPN

To prevent use of ciphers greater than DES, pre-deployment checks are available at the following locations in the Cloud-Delivered Firewall Management Center:

Devices > **Platform Settings**, click **Edit** and then **SSL**.

Devices > **VPN** > **Remote Access**, click **Edit**, and then from **Advanced** drop-down list, choose **IPsec**.

For more information about SSL settings and IPsec, see Configure SSL settings, on page 1278 and Configure Remote Access VPN IPsec/IKEv2 Parameters, on page 1866.

Authentication, Authorization, and Accounting

Configure DNS on each device in the topology in to use remote access VPN. Without DNS, the device cannot resolve AAA server names, named URLs, and CA Servers with FQDN or Hostnames; it can only resolve IP addresses.

You can configure DNS using the **Platform Settings**. For more information, see Configure DNS server settings, on page 1246 and DNS server groups, on page 2965.

Client Certificates

If you are using client certificates in your deployment, they must be added to your client's platform independent of the Secure Firewall Threat Defense or Cloud-Delivered Firewall Management Center. Facilities such as SCEP or CA Services are not provided to populate your clients with certificates.

Unsupported Features of Secure Client

The only supported VPN client is the Cisco Secure Client. No other clients or native VPNs are supported. Clientless VPN is not supported for VPN connectivity; it is only used to deploy the Secure Client using a web browser.

Note Using multiple Secure Client packages on Firewall Threat Defense devices can increase memory usage and affect the device's performance. We recommend that you do not use multiple Secure Client packages on low-end threat defense devices to avoid continuous reboots because of memory exhaustion.

The following Secure Client features are not supported when connecting to a Firewall Threat Defense secure gateway:

- TACACS, Kerberos (KCD Authentication and RSA SDI), and SDI.

Configuring a New Remote Access VPN Connection

This section provides instructions to configure a new remote access VPN policy with Secure Firewall Threat Defense devices as VPN gateways and Cisco Secure Client as the VPN client.

Step	Do This	More Info
1	Review the guidelines and prerequisites.	Guidelines and Limitations for Remote Access VPNs, on page 1818 Prerequisites for Configuring Remote Access VPN, on page 1822
2	Create a new remote access VPN policy using the wizard.	Create a New Remote Access VPN Policy, on page 1822
3	Update the access control policy deployed on the device.	Update the Access Control Policy on the Secure Firewall Threat Defense Device, on page 1825
4	(Optional) Configure a NAT exemption rule if NAT is configured on the device.	(Optional) Configure NAT Exemption, on page 1826
5	Configure DNS.	Configure DNS, on page 1827
6	Add Secure Client Profile.	Add Secure Client Profile XML File, on page 1827
7	Deploy the remote access VPN policy.	Deploy configuration changes, on page 202
8	(Optional) Verify the remote access VPN policy configuration.	Verify the Configuration, on page 1830

Prerequisites for Configuring Remote Access VPN

- Deploy Secure Firewall Threat Defense devices and configure Cloud-Delivered Firewall Management Center to manage the device with required licenses with export-controlled features enabled. For more information, see VPN licensing, on page 1714.
- Configure the certificate enrollment object that is used to obtain the identity certificate for each Firewall Threat Defense device that act as a remote access VPN gateway.
- Configure any AD or LDAP realms being used by remote access VPN policies.
- During migration of Firewall Threat Defense with remote access VPN, the realm (LDAP, AD or even local) used in remote access VPN policy should be preconfigured on Cloud-Delivered Firewall Management Center before you migrate the policy.
- Ensure that the AAA server is reachable from the Firewall Threat Defense device. Configure routing (**Devices** > **Device Management**, click the **Edit** icon. From the left pane, choose **Routing**) to ensure connectivity to the AAA servers.

  For remote access VPN double authentication, ensure that both the primary and secondary authentication servers are reachable from the Firewall Threat Defense device.
- Purchase and enable one of the following Cisco Secure Client licenses: Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only to enable the Firewall Threat Defense remote access VPN.
- Download the latest Secure Client image files from Cisco Software Download Center.

  On your Cloud-Delivered Firewall Management Center web interface, go to **Objects** > **Object Management** > **VPN** > **Secure Client File** and add the new Secure Client image files.
- Create a security zone or interface group that contains the network interfaces that users will access for VPN connections. See Interfaces, on page 2980.
- Download the Secure Client Profile Editor from Cisco Software Download Center to create the Secure Client profile. You can use the standalone profile editor to create a new or modify an existing Secure Client profile.

Create a New Remote Access VPN Policy

The Remote Access VPN Policy Wizard guides you to quickly and easily set up remote access VPNs with basic capabilities. You can further enhance the policy configuration by specifying additional attributes as you want and deploy it to your Secure Firewall Threat Defense secure gateway devices.

Before you begin

- Ensure that you complete all the prerequisites listed in Prerequisites for Configuring Remote Access VPN, on page 1822.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Add** to create a new remote access VPN policy with basic policy configuration, using the Remote Access VPN Policy wizard.

You must proceed through the entire wizard to create a new policy; the policy is not saved if you cancel before you complete the wizard.

Step 3 Select the target devices and protocols.

The Firewall Threat Defense devices that you select here functions as your remote access VPN gateways for the VPN client users.

You can select Firewall Threat Defense devices when you create a remote access VPN policy or change them later. See Set Target Devices for a Remote Access VPN Policy, on page 1831.

You can select **SSL** or **IPSec-IKEv2**, or both the VPN protocols. Firewall Threat Defense supports both the protocols to establish secure connections over a public network through VPN tunnels.

Note
Firewall Threat Defense does not support IPSec tunnels with NULL encryption. If you have selected IPSec-IKEv2, make sure that you do not choose NULL encryption for IPSec IKEv2 proposal. See Configure an IKEv2 IPsec proposal object, on page 3071.

For SSL settings, see Configure SSL settings, on page 1278.

Step 4 Click **Next**.

Step 5 Configure the **Connection Profile** and **Group Policy** settings.

A connection profile specifies a set of parameters that define how the remote users connect to the VPN device. The parameters include settings and attributes for authentication, address assignments to VPN clients, and group policies. Firewall Threat Defense device provides a default connection profile named *DefaultWEBVPNGroup* when you configure a remote access VPN policy.

For more information, see Configure Connection Profile Settings, on page 1832.

Step 6 Configure the **Authentication, Authorization & Accounting** settings.

For information about configuring,

- AAA settings, see Configure AAA Settings for Remote Access VPN, on page 1835
- LDAP attribute maps, see Configuring LDAP Attribute Mapping, on page 1858
- SAML 2.0 single sign-on authentication, see Configuring a SAML Single Sign-On Authentication, on page 1908

Step 7 Configure the **Client Address Assignment** settings.

Client IP address can be assigned from AAA server, DHCP server and IP address pools. When multiple options are selected, IP address assignment is done in the order of AAA server, DHCP server, and IP address pool. Assignment of client IP addresses from the AAA server is supported only for realm and RADIUS authorization. Ensure that realm or RADIUS server is configured to provide client IP address.

Step 8 Configure the **Group Policy** settings.

A group policy is a set of attribute and value pairs, stored in a group policy object, that define the remote access VPN experience for VPN users. You configure attributes such as user authorization profile, IP addresses, Secure Client settings, VLAN mapping, and user session settings and so on using the group policy. The

RADIUS authorization server assigns the group policy, or it is obtained from the current connection profile. For more information, see Configuring Group Policies, on page 1857.

Step 9 Click **Next**.

Step 10 Select the **Secure Client Image** that the VPN users will use to connect to the remote access VPN.

The Secure Client provides secure SSL or IPSec (IKEv2) connections to the Secure Firewall Threat Defense device for remote users with full VPN profiling to corporate resources. After the remote access VPN policy is deployed on the Firewall Threat Defense device, VPN users can enter the IP address of the configured device interface in their browser to download and install the Secure Client.

For information about configuring the client profile and client modules, see Group policy options for Secure Client, on page 3065.

Step 11 Click **Next**.

Step 12 Configure **Network Interface for Incoming VPN Access**.

Interface objects segment your network to help you manage and classify traffic flow. A security zone object simply groups interfaces. These groups may span multiple devices; you can also configure multiple zones interface objects on a single device. There are two types of interface objects:

- Security zones—An interface can belong to only one security zone.
- Interface groups—An interface can belong to multiple interface groups (and to one security zone).

(Optional) Check the **Enable DTLS on member interfaces** check box, if required. DTLS is applicable only for SSL protocol.

Step 13 Configure **Device Certificates**.

Device certificate (also called Identity certificate) identifies the VPN gateway to the remote access clients. Select a certificate which is used to authenticate the VPN gateway. From the **Certificate Enrollment** drop-down list, choose a certificate or click **+** to add a certificate to authenticate the VPN gateway.

Step 14 Configure **Service Access Control**.

You can use a service access object to control remote clients' access to VPN on Firewall Threat Defense devices with Version 7.7 or later. This object provides geolocation-based access control to clients before VPN authentication. By default, there is no access control for RAVPN, and remote clients can connect from any geolocation unless specified by a service access object. For more information, see Manage VPN Access of Remote Users Based on Geolocation, on page 1886and Configure a service access object, on page 2952.

Step 15 Configure **Access Control for VPN Traffic** .

By default, all decrypted traffic in the VPN tunnel is subjected to the Access Control Policy. Check the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** check box to bypass decrypted traffic from the Access Control Policy. This option bypasses the Access Control Policy inspection, but VPN filter ACL and authorization ACL downloaded from AAA server are still applied to VPN traffic.

Note
If you select this option, you need not update the access control policy for remote access VPN as specified in Update the Access Control Policy on the Secure Firewall Threat Defense Device, on page 1825.

Step 16 Click **Next**.

Step 17 View the **Summary** of the remote access VPN policy configuration.

The Summary page displays all the remote access VPN settings you have configured so far and provides links to the additional configurations that need to be performed before deploying the remote access VPN policy on the selected devices.

Click **Back** to make changes to the configuration, if required.

Step 18 Click **Finish** to complete the basic configuration for the remote access VPN policy.

When you complete the Remote Access VPN Policy Wizard, the policy listing page appears. Later, set up DNS configuration, configure access control for VPN users, and enable NAT exemption (if necessary) to complete a basic remote access VPN Policy configuration.

What to do next

Use the Remote Access VPN dashboard (**Overview** > **Dashboards** > **Remote Access VPN**) to monitor real-time data from active remote access VPN sessions on the devices. You can quickly determine problems related to user sessions and mitigate the problems for your network and users.

Update the Access Control Policy on the Secure Firewall Threat Defense Device

Before deploying the remote access VPN policy, you must update the access control policy on the targeted Secure Firewall Threat Defense device with a rule that allows VPN traffic. The rule must allow all traffic coming in from the outside interface, with source as the defined VPN pool networks and destination as the corporate network.

Note If you have selected the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option on the Access Interface tab, you need not update the access control policy for remote access VPN.

Enable or disable the option for all your VPN connections. If you disable this option, make sure that the traffic is allowed by the access control policy or pre-filter policy.

For more information, see Configure Access Interfaces for Remote Access VPN, on page 1851.

Before you begin

Complete the remote access VPN policy configuration using the Remote Access VPN Policy wizard.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** on the access control policy that you want to update.

Step 3 Click **Add Rule** to add a new rule.

Step 4 Specify the **Name** for the rule and select **Enabled**.

Step 5 Select the **Action**, **Allow** or **Trust**.

Step 6 Select the following on the **Zones** tab:

a) Select the outside zone from Available Zones and click **Add to Source**.

b) Select the inside zone from Available Zones and click **Add to Destination**.

Step 7 Select the following on the **Networks** tab:

a) Select the inside network (inside interface and/or a corporate network) from Available networks and click **Add to Destination**.

b) Select the VPN address pool network from **Available Networks** and click **Add to Source Networks**.

Step 8 Configure other required access control rule settings and click **Add**.

Step 9 Save the rule and access control policy.

(Optional) Configure NAT Exemption

NAT exemption exempts addresses from translation and allows both translated and remote hosts to initiate connections with your protected hosts. Like identity NAT, you do not limit translation for a host on specific interfaces; you must use NAT exemption for connections through all interfaces. However, NAT exemption enables you to specify the real and destination addresses when determining the real addresses to translate (similar to policy NAT). Use static identity NAT to consider ports in the access list.

When you configure static identity NAT for remote access or site-to-site VPN, you must configure NAT with the route lookup option. Without route lookup, the Firewall Threat Defense sends traffic out of the interface specified in the NAT command, regardless of what the routing table says. For example, you do not want the Firewall Threat Defense to send the DHCP scope traffic through an incorrect interface; it will never return to the interface IP address. The route lookup option lets the Firewall Threat Defense send, or intercept, the traffic directly on the interface IP address instead of through the interface. For traffic from the VPN client to a host on the inside network, the route lookup option will still result in the correct egress interface (inside), so normal traffic flow is not affected.

Before you begin

Check if NAT is configured on the targeted devices where remote access VPN policy is deployed. If NAT is enabled on the targeted devices, you must define a NAT policy to exempt VPN traffic.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, click **Devices** > **NAT**.

Step 2 Select a NAT policy to update or click **New Policy** > **Threat Defense NAT** to create a NAT policy with a NAT rule to allow connections through all interfaces.

Step 3 Click **Add Rule** to add a NAT rule.

Step 4 On the Add NAT Rule window, select the following:

a) Select the NAT Rule as **Manual NAT Rule**.

b) Select the Type as **Static**.

c) Click **Interface Objects** and select the Source and destination interface objects.

Note

This interface object must be the same as the interface selected in the remote access VPN policy.

For more information, see Configure Access Interfaces for Remote Access VPN, on page 1851.

a) Click **Translation** and select the source and destination networks:

- **Original Source** and **Translated Source**
- **Original Destination** and **Translated Destination**

Step 5 On the Advanced tab, select **Do not proxy ARP on Destination Interface**.

Do not proxy ARP on Destination Interface—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router.

Step 6 Click **OK**.

Configure DNS

Configure DNS on each Firewall Threat Defense device in order to use remote access VPN. Without DNS, the devices cannot resolve AAA server names, named URLs, and CA Servers with FQDN or Hostnames. It can only resolve IP addresses.

Procedure

Step 1 Configure DNS server details and domain-lookup interfaces using the Platform Settings. For more information, see Configure DNS server settings, on page 1246 and DNS server groups, on page 2965.

Step 2 Configure split-tunnel in group policy to allow DNS traffic through remote access VPN tunnel if the DNS server is reachable through VNP network. For more information, see Configure a group policy object, on page 3062.

Add Secure Client Profile XML File

The Secure Client Profile is a group of configuration parameters stored in an XML file that the client uses to configure its operation and appearance. These parameters (XML tags) include the names and addresses of host computers and settings to enable more client features.

You can create the Secure Client Profile using the Secure Client Profile Editor, a GUI-based configuration tool that is available as part of the Secure Client software package. It is an independent program that you run outside of the Cloud-Delivered Firewall Management Center. For more information about Secure Client Profile Editor, see Cisco Secure Client (including AnyConnect) Administrator Guide.

Before you begin

A Secure Firewall Threat Defense remote access VPN policy requires assignment of the Secure Client Profile to the VPN clients. You can attach the client profile to a group policy.

Download the Secure Client Profile Editor from Cisco Software Download Center.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy which you want to update.

Step 3 Click **Edit** on the connection profile to which you want to add the Secure Client profile.

Step 4 Click **Edit Group Policy**. If you choose to add a new group policy, click **Add**.

Step 5 Choose **Secure Client** > **Profile**.

Step 6 Choose a profile from the **Client Profile** drop-down list. If you choose to add a new client profile, click **Add** and do the following:

a) Specify the profile **Name**.

b) Click **Browse** and select the Secure Client Profile XML file.

Note
For two-factor authentication, make sure that the timeout is set to 60 seconds or more in the Secure Client profile.

c) Click **Save**.

Step 7 Save your changes.

(Optional) Configure Split Tunneling

Split tunnel allows VPN connectivity to a remote network across a secure tunnel and also to a network outside VPN tunnel. Configure split tunneling if you want to allow your VPN users to access an outside network while they remain connected to the remote access VPN. To configure a split-tunnel list, you must create a Standard Access List or Extended Access List.

For more information, see Configuring Group Policies, on page 1857.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy for which you want to configure split tunneling.

Step 3 Click **Edit** on the required connection profile.

Step 4 Click **Add** to add a group policy or click **Edit Group Policy**.

Step 5 Choose **General** > **Split Tunneling**.

Step 6 From the **IPv4 Split Tunneling** or **IPv6 Split Tunneling** list, select **Exclude networks specified below** and then select the networks that you want to exclude from VPN traffic.

The default setting allows all traffic over the VPN tunnel.

Step 7 Click **Standard Access List** or **Extended Access List**, and select an access list from the drop-down or add a new one.

Step 8 If you choose to add a new standard or extended access list, do the following:

a) Specify the **Name** for the new access list and click **Add**.
b) Choose **Allow** from the **Action** drop-down.
c) Select the network traffic that you want to allow over the VPN tunnel and click **Add**.

Step 9 Save your changes.

Related Topics

Access lists, on page 2948

(Optional) Configure Dynamic Split Tunneling

Dynamic split tunneling allows you to fine-tune split tunneling based on DNS domain names. You can configure domains that must be included or excluded in the remote access VPN tunnel. Excluded domains are not blocked. Instead, traffic to those domains is kept outside the VPN tunnel. For example, you could send traffic to Cisco WebEx on the public Internet, thus freeing bandwidth in your VPN tunnel for traffic that is targeted to servers within your protected network. For more information about configuring this feature, see Configure AnyConnect Dynamic Split Tunnel on FTD Managed by FMC.

Before you begin

You can configure this feature using the Cloud-Delivered Firewall Management Center and Firewall Threat Defense from versions 7.0 or later. If you have an older version of the Cloud-Delivered Firewall Management Center, you can configure it using FlexConfig as instructed in the Advanced AnyConnect VPN Deployments for Firepower Threat Defense with FMC.

Procedure

Step 1 Configure the group policy to use Dynamic Split Tunnel.

a) Choose **Devices** > **VPN** > **Remote Access**.
b) Click **Edit** on the remote access VPN policy for which you want to configure dynamic split tunneling.
c) Click **Edit** on the required connection profile.
d) Click **Edit Group Policy**.

Step 2 Configure the Secure Client custom attribute in the **Add/Edit Group Policy** dialog box.

a) Click the Secure Client tab.
b) Click **Custom Attributes** and click **+**.
c) Choose **Dynamic Split Tunneling** from the **Secure Client Attribute** drop-down list.
d) Click **+** to create a new custom attribute object.
e) Enter the name for the custom attribute object.
f) **Include domains**—Specify domain names that will be included in the remote access VPN tunnel.

You can include domains in the tunnel that will be excluded based on IP addresses.

g) **Exclude domains**—Specify domain names that will be excluded from the remote access VPN.

Excluded domains are not blocked, traffic to these domains is kept outside the VPN tunnel.

h) Click **Save**.

i) Click **Add**.

Step 3 Verify the configured custom attribute and click **Save** to save the group policy.

Step 4 Click **Save** to save the connection profile.

Step 5 Click **Save** to save the remote access VPN policy.

What to do next

1. Deploy the configuration to Firewall Threat Defense.
2. Verify the configured dynamic split tunnel configuration on the Firewall Threat Defense and the Secure Client. For more information, see Verify Dynamic Split Tunneling Configuration, on page 1830.

Verify Dynamic Split Tunneling Configuration

On the Firewall Threat Defense

Use the following commands to verify the dynamic split tunneling configuration:

- **show running-config webvpn**
- **show running-config anyconnect-custom-data**
- **show running-config group-policy** <*group-policy-name*>

On the Secure Client

Click the Statistics () icon and choose **VPN > Statistics**. You can confirm the domains under the Dynamic Split Exclusion/Inclusion category.

Verify the Configuration

Procedure

Step 1 Open a web browser on a machine on the outside network.

Step 2 Enter the URL of the Firewall Threat Defense remote access VPN gateway device.

Step 3 Enter the username and password when prompted, and click **Logon**.

Note
Connection to the VPN establishes automatically if you install Secure Client on the system.

The VPN prompts you to download Secure Client if Secure Client is not installed.

Step 4 Download Secure Client if it is not installed and connect to the VPN.
The Secure Client installs itself. On successful authentication, you establish connection to the Secure Firewall Threat Defense remote access VPN gateway. The remote access VPN enforces the applicable identity or QoS policy according to your VPN policy configuration.

Create a Copy of an Existing Remote Access VPN Policy

You can copy an existing remote access VPN policy to create a new one with all the settings, including the connection profiles and access interfaces. You can then assign devices to the new policy and deploy the VPN on the assigned devices as required.

Note Users with read-only permission for remote access VPN cannot create copy of the VPN. Users with read-only privileges in the domain can copy the remote access VPNs.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.
Step 2 Click **Copy** on the policy that you want to copy.
Step 3 Specify a **Name** for the new remote access VPN.
Step 4 Click **OK**.

What to do next

To assign devices to the new policy, see Set Target Devices for a Remote Access VPN Policy, on page 1831.

Set Target Devices for a Remote Access VPN Policy

After you create remote access VPN policy, you can assign the policy to threat defense devices.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.
Step 2 Click **Edit** ( ) next to the remote access VPN policy that you want to edit.
Step 3 Click **Policy Assignments**.
Step 4 Do any of the following:

- To assign a device, high-availability pair, or device group to the policy, select it in the **Available Devices** list and click **Add**. You can also drag and drop the available devices to select.

- To remove a device assignment, click **Delete** ( ) next to a device, high-availability pair, or device group in the **Selected Devices** list.

Step 5 Click **OK**.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes.

Associate Local Realm with Remote Access VPN Policy

You can associate local realm to remote access VPN policy to enable local user authentication.

For information about creating and managing realms, see Manage a realm, on page 2341.

For information about configuring local user authentication for remote access VPNs, see Configure AAA Settings for Remote Access VPN, on page 1835.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** ( ) next to the remote access VPN policy that you want to edit.

Step 3 Click the link next to **Local Realm**.

Step 4 Select the **Local Realm Server** from the list, or click **Add** to add a new local realm.

Step 5 Click **OK**.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes.

Additional Remote Access VPN Configurations

Configure Connection Profile Settings

Remote Access VPN policy contains the connection profiles targeted for specific devices. These policies pertain to creating the tunnel itself, such as, how AAA is accomplished, and how addresses are assigned (DHCP or Address Pools) to VPN clients. They also include user attributes, which are identified in group policies configured on the Firewall Threat Defense device or obtained from a AAA server. A device also

provides a default connection profile named *DefaultWEBVPNGroup*. The connection profile that is configured using the wizard appears in the list.

If you decide to grant different rights to different groups of VPN users, then you can add specific connection profiles for each of the user groups and maintain multiple connection profiles in your remote access VPN policy.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Select a **Connection Profile** and click **Edit**.

Step 4 (Optional) If you choose to add new connection profile, click **Add**.

Step 5 Configure IP Addresses for VPN Clients.
Configure IP Addresses for VPN Clients, on page 1833

Step 6 (Optional) Update AAA Settings for remote access VPNs.
Configure AAA Settings for Remote Access VPN, on page 1835

Step 7 (Optional) Create or update Aliases.
Create or Update Aliases for a Connection Profile, on page 1850

Step 8 Save your changes.

Configure IP Addresses for VPN Clients

Client address assignment allows you to assign IP addresses for the remote access VPN users.

You can assign IP Address for remote VPN clients from the local IP address pools, DHCP Servers, and AAA servers. The AAA servers are assigned first, followed by others. Configure the **Client Address Assignment** policy in the **Advanced** tab to define the assignment criteria. The IP pools defined in this connection profile will only be used if no IP pools are defined in group policy associated with the connection profile, or the system default group policy **DfltGrpPolicy**.

IPv4 Address Pools—SSL VPN clients receive new IP addresses when they connect to the Firewall Threat Defense device. Address pools define a range of addresses that remote clients can receive. You can add a maximum of six pools for IPv4 and IPv6 addresses each.

Note You can use the IP address from the existing IP pools in the Cloud-Delivered Firewall Management Center or create a new pool using the **Add** option. Also, you can create an IP pool in Cloud-Delivered Firewall Management Center using the **Objects** > **Object Management** > **Address Pools** path. For more information, see Configure address pools, on page 2954.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Existing remote access policies are listed.

Step 2 Select a remote access VPN policy and click the edit icon.

Step 3 Select the connection profile that you want to update and click the edit icon.

Step 4 Under the **Client Address Assignment** tab, do the following:

Step 5 Click + next to **Address Pools**:

a) Click + next to **Address Pools** to add IP addresses, and select **IPv4** or **IPv6** to add the corresponding address pool. Select the IP address pool from **Available Pools** and click **Add**.

Note
If you share your remote access VPN policy among multiple Secure Firewall Threat Defense devices, bear in mind that all devices share the same address pool unless you use device-level object overrides to replace the global definition with a unique address pool for each device. Unique address pools are required to avoid overlapping addresses in cases where the devices are not using NAT.

b) Click + next to **Available Pools** in the **Address Pools** window to add a new IPv4 or IPv6 address pool. When you choose the IPv4 pool, provide a starting and ending IP address. When you choose to include a new IPv6 address pool, enter **`Number of Addresses`** in the range 1-16384. Select the **Allow Overrides** option to avoid conflicts with IP address when objects are shared across many devices. For more information, see Configure address pools, on page 2954.

c) Click **OK**.

If you plan to edit the IP address pools, we recommend that you perform the following steps during a maintenance window:

1. Unassign the device from the remote access VPN.
2. Select the device and click **Deploy**.

   This deployment removes all the remote access VPN configurations from the device, terminates the remote access VPN sessions, the sessions are not reestablished.
3. Click the edit icon next to the IP address pools to edit it, edit any other remote access VPN configurations, if required, on the Cloud-Delivered Firewall Management Center.
4. Assign the device to the updated remote access VPN policy.
5. Deploy the configurations on the device.

   The remote access VPN clients can connect to the device after the maintenance window.

Step 6 Click + next to **DHCP Servers** to add DHCP servers:

Note
The DHCP server address can be configured only with IPv4 address.

a) Specify the name and DHCP (Dynamic Host Configuration Protocol) server address as network objects. Click **Add** to choose the server from the object list. Click **Delete** to delete a DHCP server.

b) Click **Add** in the **New Objects** page to add a new network object. Enter the new object name, description, network, and select the **Allow Overrides** option as applicable. For more information, see Create a network object, on page 2985 and Allow object overrides, on page 2941.

c) Click **OK**.

Step 7 Click **Save**.

Related Topics

Configure Connection Profile Settings, on page 1832

Configure AAA Settings for Remote Access VPN

Before you begin

- Ensure that the required machine and user certificates are deployed on the endpoints. For information about Firewall Threat Defense certificates, see Manage Firewall Threat Defense certificates, on page 3080.
- Configure Secure Client profiles with required certificates. For more information, see *Cisco Secure Client (including AnyConnect) Administrator Guide*.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Select a connection profile to update AAA settings, click **Edit** and then choose**AAA**.

Step 4 Select the following for **Authentication**:

- **Authentication Method**—Determines how a user is identified before being allowed access to the network and network services. It controls access by requiring valid user credentials, which are typically a username and password. It may also include the certificate from the client. Supported authentication methods are **AAA only**, **Client Certificate only**, and **AAA + Client Certificate**.

  When you select the **Authentication Method** as:

  - **AAA Only**—If you select the **Authentication Server** as **RADIUS**, by default, the Authorization Server has the same value. Select the **Accounting Server** from the drop-down list. Whenever you select **AD** and **LDAP** from the Authentication Server drop-down list, you must select the **Authorization Server** and **Accounting Server** manually.
  - **SAML**—Each user is authenticated using the SAML single sign-on server. For more information, see Single Sign-On Authentication with SAML 2.0, on page 1906.

    Override Identity Provider Certificate—Select to override the primary identity provider certificate of the SAML provider with an IdP certificate specific to a connection profile or SAML application. Select the IdP certificate from the drop-down.

    Microsoft Azure can support multiple applications for the same entity ID. Each application (mapped to a different connection profile) requires a unique certificate. If you want to retain an existing entity ID for the single-sign-on object in current connection profile and use a different IdP certificate, you can select this option.

    This enables support for multiple SAML applications per Microsoft Azure SAML identity provider.

    The primary identity certificate is configured in the single sign-on server object.

    For information about configuring a single sign-on server object, see Add a single sign-on server, on page 2946.

    Choose your **SAML Login Experience** to configure a browser for SAML web authentication:

- **VPN client embedded browser**—Choose this option to use the browser embedded with the VPN client for web authentication. The authentication applies to the VPN connection only.
- **Default OS Browser**—Choose this option to configure the operating system that default or native browser that supports WebAuthN (FIDO2 standard for web authentication). This option enables single sign-on(SSO) support for web authentication methods such as biometric authentication.

  The default browser requires an external browser package for web authentication. The package Default-External-Browser-Package is configured by default. You can change the default external browser package by editing a remote access VPN policy and selecting the file under **Advanced** and choose **AnyConnect Client ImagesSecure Client Image**s, then click **Package File**.

  You can also add a new package file by selecting. **Objects** > **Object Management** > **VPN** > **Secure Client File** and then click **Add Secure Client File**.

- **Client Certificate Only**—Each user is authenticated with a client certificate. The client certificate must be configured on VPN client endpoints. By default, the user name is derived from the client certificate fields CN and OU. If the user name is specified in other fields in the client certificate, use 'Primary' and 'Secondary' field to map appropriate fields.

  Select **Enable multiple certificate authentication** to authenticate the VPN client using the machine and user certificates.

  If have enabled multiple certificate authentication, you can select one of the following certificates to map the username and authenticate the VPN user:

  - **First Certificate**—Select this option to map the username from the machine certificate sent from the VPN client.
  - **Second Certificate**—Select this option to map the username from the user certificate sent from the client.

  Note
  If you do not enable multiple certificate authentication, the user certificate (second certificate) is used for authentication by default.

  If you select the **Map specific field** option, which includes the username from the client certificate, the **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively. If you select the **Use entire DN as username** option, the system automatically retrieves the user identity. A distinguished name (DN) is a unique identification, made up of individual fields used as the identifier when matching users to a connection profile. DN rules are used for enhanced certificate authentication.

  The primary and Secondary fields pertaining to the **Map specific field** option contain these common values:

  - C (Country)
  - CN (Common Name)
  - DNQ (DN Qualifier)
  - EA (Email Address)
  - GENQ (Generational Qualifier)

- GN (Given Name)
- I (Initial)
- L (Locality)
- N (Name)
- O (Organisation)
- OU (Organisational Unit)
- SER (Serial Number)
- SN (Surname)
- SP (State Province)
- T (Title)
- UID (User ID)
- UPN (User Principal Name)

- **Client Certificate & AAA**— Each user is authenticated with both a client certificate and AAA server. Select the required certificate and AAA configurations for authentication.

  Whichever authentication method you choose, select or deselect **Allow connection only if user exists in authorization database**.

- **Client Certificate & SAML**— Each user is authenticated with both a client certificate and SAML server. Select the required certificate and SAML configurations for authentication.
  - **Allow connection only if username from certificate and SAML are the same**—Select to allow VPN connection only if the username from the certificate matches the SAML single sign-on username.
  - **Use username from client certificate for Authorization**—When you choose the option to pick username from client certificate for authorization, you must configure the fields to pick from the client certificate.

    You can choose to map a specific filed as the username or use the entire distinguished name (DN) for authorization:
    - **Map specific field**— Select to include the username from the client certificate; the **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively.
    - **Use entire DN as username**— The system automatically retrieves the user identity for authorization.

You can also create a Dynamic Access Policy (DAP) to match user-specific SAML assertion attributes or username to DAP certificate attributes. See Configure AAA criteria settings for a DAP record, on page 1932.

- **Authentication Server**—Authentication is the way a user is identified before being allowed access to the network and network services. Authentication requires valid user credentials, a certificate, or both. You can use authentication alone, or with authorization and accounting.

  Select an authentication server from the list if you have added a server already or else create an authentication server:

  - **LOCAL**—Use a local database from the Firewall Threat Defense for user authentication. To configure local authentication, Firewall Threat Defense must be Version 7.0 and later.
    - **Local Realm**—Select a local realm or click **Add** to configure a realm. See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
  - **Realm**—Configure an LDAP or AD realm. See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
  - **RADIUS Server Group**—Add a RADIUS server group object with RADIUS servers. See Add a RADIUS server group, on page 2943.
  - **Single Sign-On Server**—Create a single sign-on server object for SAML authentication. See Add a single sign-on server, on page 2946.

Fallback to LOCAL Authentication— The user is authenticated using the local database and the VPN tunnel can be established even if the AAA server group is unavailable, provided that the local database is configured.

- **Use secondary authentication** — Secondary authentication is configured in addition to primary authentication to provide additional security for VPN sessions. Secondary authentication is applicable only to **AAA only** and **Client Certificate & AAA** authentication methods.

  Secondary authentication is an optional feature that requires a VPN user to enter two sets of username and password on the Secure Client login screen. You can also configure to pre-fill the secondary username from the authentication server or client certificate. Remote access VPN authentication is granted only if both primary and secondary authentications are successful. VPN authentication is denied if any one of the authentication servers is not reachable or one authentication fails.

  You must configure a secondary authentication server group (AAA server) for the second username and password before configuring secondary authentication. For example, you can set the primary authentication server to an LDAP or Active Directory realm and the secondary authentication to a RADIUS server.

  Note
  By default, secondary authentication is not required.

  Authentication Server— Secondary authentication server to provide secondary username and password for VPN users.

  - **Fallback to LOCAL Authentication**— This user is authenticated using the local database and the VPN tunnel can be established even if the AAA server group is unavailable, provided that the local database is configured.

  Select the following under **Username for secondary authentication**:

  - **Prompt**: Prompts the users to enter the username and password while logging on to VPN gateway.
  - **Use primary authentication username**: The username is taken from the primary authentication server for both primary and secondary authentication; you must enter two passwords.
  - **Map username from client certificate**: Prefills the secondary username from the client certificate.

If you have enabled multiple certificate authentication, you can select one of the following certificates:

- **First Certificate**— Select this option to map the username from the machine certificate sent from the VPN client.
- **Second Certificate**— Select this option to map the username from the user certificate sent from the client.
- If you select **Map specific field** option, which includes the username from the client certificate. The **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively. If you select the **Use entire DN (Distinguished Name) as username** option, the system automatically retrieves the user identity.

  See **Authentication Method** descriptions for more information about primary and secondary field mapping.
- **Prefill username from certificate on user login window**: Prefills the secondary username from the client certificate when the user connects via Secure Client.
  - **Hide username in login window**: The secondary username is pre-filled from the client certificate, but hidden to the user so that the user does not modify the pre-filled username.

- **Use secondary username for VPN session**: The secondary username is used for reporting user activity during a VPN session.

Step 5 Select the following for **Authorization**:

- **Authorization Server**—After authentication is complete, authorization controls the services and commands available to each authenticated user. Authorization works by assembling a set of attributes that describe what the user is authorized to perform, their actual capabilities, and restrictions. When you do not use authorization, authentication alone provides the same access to all authenticated users. Authorization requires authentication.

  To know more about how remote access VPN authorization works, see Understanding Policy Enforcement of Permissions and Attributes, on page 1816.

  When a RADIUS Server is configured for user authorization in the connection profile, the remote access VPN system administrator can configure multiple authorization attributes for users or user-groups. The authorization attributes that are configured on the RADIUS server can be specific for a user or a user-group. Once users are authenticated, these specific authorization attributes are pushed to the Firewall Threat Defense device.

  Note
  The AAA server attributes obtained from the authorization server override the attribute values that may have been previously configured on the group policy or the connection profile.

- Check **Allow connection only if user exists in authorization database** if desired.

  When enabled, the system checks the username of the client must exist in the authorization database to allow a successful connection. If the username does not exist in the authorization database, then the connection is denied.

- When you select a realm as the Authorization Server, you must configure an LDAP attribute map. You can choose a single server for authentication and authorization or a different servers. Click **Configure LDAP Attribute Map** to add LDAP attribute maps for authorization.

Note
Firewall Threat Defense does not support SAML Identity Provider as the Authorization server. If the Active Directory behind the SAML Identity Provider is reachable via Cloud-Delivered Firewall Management Center and Firewall Threat Defense, you can configure authorization following these steps:

- Add realm for the AD Server. See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
- Select the realm object as the Authorization Server in remote access VPN connection profile.
- Configure LDAP attribute map for the selected realm.

For information about configuring LDAP attribute maps, see Configuring LDAP Attribute Mapping, on page 1858.

Step 6 Select the following for **Accounting**:

- **Accounting Server**—Accounting is used to track the services that users are accessing and the amount of network resources they are consuming. When AAA accounting is activated, the network access server reports user activity to the RADIUS server. Accounting information includes when sessions start and stop, usernames, the number of bytes that pass through the device for each session, the services used, and the duration of each session. This data can then be analyzed for network management, client billing, or auditing. You can use accounting alone or together with authentication and authorization.

  Specify the RADIUS Server Group object that will be used to account for the remote access VPN session.

Step 7 Select the following **Advanced Settings**:

- **Strip Realm from username**—Select to remove the realm from the username before passing the username on to the AAA server. For example, if you select this option and provide *domain\username*, the domain is stripped off from the username and sent to AAA server for authentication. By default this option is unchecked.
- **Strip Group from username**—Select to remove the group name from the username before passing the username on to the AAA server. By default this option is unchecked.

  Note
  A realm is an administrative domain. Enabling these options allows the authentication to be based on the username alone. You can enable any combination of these options. However, you must select both check boxes if your server cannot parse delimiters.

- **Password Management**—Enable managing the password for the remote access VPN users. Select to notify ahead of the password expiry or on the day the password expires.

Step 8 Click **Save**.

Related Topics

Understanding Policy Enforcement of Permissions and Attributes, on page 1816
Manage a realm, on page 2341

RADIUS Server Attributes for Secure Firewall Threat Defense

The Firewall Threat Defense device supports applying user authorization attributes (also called user entitlements or permissions) to VPN connections from the external RADIUS server that are configured for authentication and/or authorization in the remote access VPN policy.

Note Secure Firewall Threat Defense devices support attributes with vendor ID 3076.

The following user authorization attributes are sent to the Firewall Threat Defense device from the RADIUS server.

- RADIUS attributes 146 and 150 are sent from Firewall Threat Defense devices to the RADIUS server for authentication and authorization requests.
- All three (146, 150, and 151) attributes are sent from Firewall Threat Defense devices to the RADIUS server for accounting start, interim-update, and stop requests.

Table 104: RADIUS Attributes Sent from Secure Firewall Threat Defense to RADIUS Server

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Connection Profile Name or Tunnel Group Name	146	String	Single	1-253 characters
Client Type	150	Integer	Single	2 = Secure Client SSL VPN, 6 = Secure Client IPsec VPN (IKEv2)
Session Type	151	Integer	Single	1 = Secure Client SSL VPN, 2 = Secure Client IPsec VPN (IKEv2)

Table 105: Supported RADIUS Authorization Attributes

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Access-Hours	Y	1	String	Single	Name of the time range, for example, Busine
Access-List-Inbound	Y	86	String	Single	Both of the Access-List attributes take the na ACL that is configured on the Firewall Threa device. Create these ACLs using the Smart CL Access List object type (select **Device** > **Adv Configuration** > **Smart CLI** > **Objects**). These ACLs control traffic flow in the inbou entering the Firewall Threat Defense device) (traffic leaving the Firewall Threat Defense direction.
Access-List-Outbound	Y	87	String	Single	

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Address-Pools	Y	217	String	Single	The name of a network object defined on the Fi Threat Defense device that identifies a subnet, wh be used as the address pool for clients connectin remote access VPN. Define the network object o **Objects** page and then associate the network obj a group policy or a connection profile.
Allow-Network-Extension-Mode	Y	64	Boolean	Single	0 = Disabled 1 = Enabled
Authenticated-User-Idle-Timeout	Y	50	Integer	Single	1-35791394 minutes
Authorization-DN-Field	Y	67	String	Single	Possible values: UID, OU, O, CN, L, SP, C, EA, GN, SN, I, GENQ, DNQ, SER, use-entire-name
Authorization-Required		66	Integer	Single	0 = No 1 = Yes
Authorization-Type	Y	65	Integer	Single	0 = None 1 = RADIUS 2 = LDAP
Banner1	Y	15	String	Single	Banner string to display for Cisco VPN remote a sessions: IPsec IKEv1, Secure Client SSL-TLS/DTLS/IKEv2, and Clientless SSL
Banner2	Y	36	String	Single	Banner string to display for Cisco VPN remote a sessions: IPsec IKEv1, Secure Client SSL-TLS/DTLS/IKEv2, and Clientless SSL. The string is concatenated to the Banner1 string , if con
Cisco-IP-Phone-Bypass	Y	51	Integer	Single	0 = Disabled 1 = Enabled
Cisco-LEAP-Bypass	Y	75	Integer	Single	0 = Disabled 1 = Enabled
Client Type	Y	150	Integer	Single	1 = Cisco VPN Client (IKEv1) 2 = Secure Clien VPN 3 = Clientless SSL VPN 4 = Cut-Through- = L2TP/IPsec SSL VPN 6 = Secure Client IPsec (IKEv2)
Client-Type-Version-Limiting	Y	77	String	Single	IPsec VPN version number string
DHCP-Network-Scope	Y	61	String	Single	IP Address
Extended-Authentication-On-Rekey	Y	122	Integer	Single	0 = Disabled 1 = Enabled
Framed-Interface-Id	Y	96	String	Single	Assigned IPv6 interface ID. Combines with Framed-IPv6-Prefix to create a complete assigne address. For example: Framed-Interface-ID=1:1 combined with Framed-IPv6-Prefix=2001:0db8::/ the assigned IP address 2001:0db8::1:1:1:1.

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Framed-IPv6-Prefix	Y	97	String	Single	Assigned IPv6 prefix and length. Combines Framed-Interface-Id to create a complete ass address. For example: prefix 2001:0db8::/64 with Framed-Interface-Id=1:1:1:1 gives the 2001:0db8::1:1:1:1. You can use this attribut an IP address without using Framed-Interfac assigning the full IPv6 address with prefix le for example, Framed-IPv6-Prefix=2001:0db
Group-Policy	Y	25	String	Single	Sets the group policy for the remote access V You can use one of the following formats: • *group policy name* • OU=*group policy name* • OU=*group policy name*;
IE-Proxy-Bypass-Local		83	Integer	Single	0 = None 1 = Local
IE-Proxy-Exception-List		82	String	Single	New line (\n) separated list of DNS domains
IE-Proxy-PAC-URL	Y	133	String	Single	PAC address string
IE-Proxy-Server		80	String	Single	IP address
IE-Proxy-Server-Policy		81	Integer	Single	1 = No Modify 2 = No Proxy 3 = Auto detec Concentrator Setting
IKE-KeepAlive-Confidence-Interval	Y	68	Integer	Single	10-300 seconds
IKE-Keepalive-Retry-Interval	Y	84	Integer	Single	2-10 seconds
IKE-Keep-Alives	Y	41	Boolean	Single	0 = Disabled 1 = Enabled
Intercept-DHCP-Configure-Msg	Y	62	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Allow-Passwd-Store	Y	16	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Authentication		13	Integer	Single	0 = None 1 = RADIUS 2 = LDAP (authoriza = NT Domain 4 = SDI 5 = Internal 6 = RAD Expiry 7 = Kerberos/Active Directory
IPsec-Auth-On-Rekey	Y	42	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Backup-Server-List	Y	60	String	Single	Server Addresses (space delimited)
IPsec-Backup-Servers	Y	59	String	Single	1 = Use Client-Configured list 2 = Disable and list 3 = Use Backup Server list

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
IPsec-Client-Firewall-Filter-Name		57	String	Single	Specifies the name of the filter to be pushed to tl as firewall policy
IPsec-Client-Firewall-Filter-Optional	Y	58	Integer	Single	0 = Required 1 = Optional
IPsec-Default-Domain	Y	28	String	Single	Specifies the single default domain name to sen client (1-255 characters).
IPsec-IKE-Peer-ID-Check	Y	40	Integer	Single	1 = Required 2 = If supported by peer certificate not check
IPsec-IP-Compression	Y	39	Integer	Single	0 = Disabled 1 = Enabled
IPsec-Mode-Config	Y	31	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Over-UDP	Y	34	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Over-UDP-Port	Y	35	Integer	Single	4001- 49151. The default is 10000.
IPsec-Required-Client-Firewall-Capability	Y	56	Integer	Single	0 = None 1 = Policy defined by remote FW Are-You-There (AYT) 2 = Policy pushed CPP 4 = from server
IPsec-Sec-Association		12	String	Single	Name of the security association
IPsec-Split-DNS-Names	Y	29	String	Single	Specifies the list of secondary domain names to the client (1-255 characters).
IPsec-Split-Tunneling-Policy	Y	55	Integer	Single	0 = No split tunneling 1 = Split tunneling 2 = Loc permitted
IPsec-Split-Tunnel-List	Y	27	String	Single	Specifies the name of the network or ACL that d the split tunnel inclusion list.
IPsec-Tunnel-Type	Y	30	Integer	Single	1 = LAN-to-LAN 2 = Remote access
IPsec-User-Group-Lock		33	Boolean	Single	0 = Disabled 1 = Enabled
IPv6-Address-Pools	Y	218	String	Single	Name of IP local pool-IPv6
IPv6-VPN-Filter	Y	219	String	Single	ACL value
L2TP-Encryption		21	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 4 = 8 = Stateless-Req 15= 40/128-Encr/Stateless-Re
L2TP-MPPC-Compression		38	Integer	Single	0 = Disabled 1 = Enabled

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Member-Of	Y	145	String	Single	Comma-delimited string, for example: `Engineering, Sales` An administrative attribute that can be used access policies. It does not set a group policy
MS-Client-Subnet-Mask	Y	63	Boolean	Single	An IP address
NAC-Default-ACL		92	String		ACL
NAC-Enable		89	Integer	Single	0 = No 1 = Yes
NAC-Revalidation-Timer		91	Integer	Single	300-86400 seconds
NAC-Settings	Y	141	String	Single	Name of the NAC policy
NAC-Status-Query-Timer		90	Integer	Single	30-1800 seconds
Perfect-Forward-Secrecy-Enable	Y	88	Boolean	Single	0 = No 1 = Yes
PPTP-Encryption		20	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 8 = Stateless-Required 15= 40/128-Encr/Stat
PPTP-MPPC-Compression		37	Integer	Single	0 = Disabled 1 = Enabled
Primary-DNS	Y	5	String	Single	An IP address
Primary-WINS	Y	7	String	Single	An IP address
Privilege-Level	Y	220	Integer	Single	An integer between 0 and 15.
Required-Client- Firewall-Vendor-Code	Y	45	Integer	Single	1 = Cisco Systems (with Cisco Integrated Cl Zone Labs 3 = NetworkICE 4 = Sygate 5 = Cis (with Cisco Intrusion Prevention Security Ag
Required-Client-Firewall-Description	Y	47	String	Single	String
Required-Client-Firewall-Product-Code	Y	46	Integer	Single	Cisco Systems Products: 1 = Cisco Intrusion Prevention Security Age Integrated Client (CIC) Zone Labs Products: 1 = Zone Alarm 2 = Zon 3 = Zone Labs Integrity NetworkICE Product: 1 = BlackIce Defende Sygate Products: 1 = Personal Firewall 2 = P Firewall Pro 3 = Security Agent
Required-Individual-User-Auth	Y	49	Integer	Single	0 = Disabled 1 = Enabled

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Require-HW-Client-Auth	Y	48	Boolean	Single	0 = Disabled 1 = Enabled
Secondary-DNS	Y	6	String	Single	An IP address
Secondary-WINS	Y	8	String	Single	An IP address
SEP-Card-Assignment		9	Integer	Single	Not used
Session Subtype	Y	152	Integer	Single	0 = None 1 = Clientless 2 = Client 3 = Client Or Session Subtype applies only when the Session T (151) attribute has the following values: 1, 2, 3,
Session Type	Y	151	Integer	Single	0 = None 1 = Secure Client SSL VPN 2 = Secur IPSec VPN (IKEv2) 3 = Clientless SSL VPN 4 = Clientless Email Proxy 5 = Cisco VPN Client (I = IKEv1 LAN-LAN 7 = IKEv2 LAN-LAN 8 = VI Balancing
Simultaneous-Logins	Y	2	Integer	Single	0-2147483647
Smart-Tunnel	Y	136	String	Single	Name of a Smart Tunnel
Smart-Tunnel-Auto	Y	138	Integer	Single	0 = Disabled 1 = Enabled 2 = AutoStart
Smart-Tunnel-Auto-Signon-Enable	Y	139	String	Single	Name of a Smart Tunnel Auto Signon list appen the domain name
Strip-Realm	Y	135	Boolean	Single	0 = Disabled 1 = Enabled
SVC-Ask	Y	131	String	Single	0 = Disabled 1 = Enabled 3 = Enable default ser Enable default clientless (2 and 4 not used)
SVC-Ask-Timeout	Y	132	Integer	Single	5-120 seconds
SVC-DPD-Interval-Client	Y	108	Integer	Single	0 = Off 5-3600 seconds
SVC-DPD-Interval-Gateway	Y	109	Integer	Single	0 = Off) 5-3600 seconds
SVC-DTLS	Y	123	Integer	Single	0 = False 1 = True
SVC-Keepalive	Y	107	Integer	Single	0 = Off 15-600 seconds
SVC-Modules	Y	127	String	Single	String (name of a module)
SVC-MTU	Y	125	Integer	Single	MTU value 256-1406 in bytes
SVC-Profiles	Y	128	String	Single	String (name of a profile)
SVC-Rekey-Time	Y	110	Integer	Single	0 = Disabled 1-10080 minutes
Tunnel Group Name	Y	146	String	Single	1-253 characters

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Tunnel-Group-Lock	Y	85	String	Single	Name of the tunnel group or “none”
Tunneling-Protocols	Y	11	Integer	Single	1 = PPTP 2 = L2TP 4 = IPSec (IKEv1) 8 = I 16 = WebVPN 32 = SVC 64 = IPsec (IKEv2) mutually exclusive. 0 - 11, 16 - 27, 32 - 43, 4 legal values.
Use-Client-Address		17	Boolean	Single	0 = Disabled 1 = Enabled
VLAN	Y	140	Integer	Single	0-4094
WebVPN-Access-List	Y	73	String	Single	Access-List name
WebVPN ACL	Y	73	String	Single	Name of a WebVPN ACL on the device
WebVPN-ActiveX-Relay	Y	137	Integer	Single	0 = Disabled Otherwise = Enabled
WebVPN-Apply-ACL	Y	102	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Auto-HTTP-Signon	Y	124	String	Single	Reserved
WebVPN-Citrix-Metaframe-Enable	Y	101	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Content-Filter-Parameters	Y	69	Integer	Single	1 = Java ActiveX 2 = Java Script 4 = Image 8 in images
WebVPN-Customization	Y	113	String	Single	Name of the customization
WebVPN-Default-Homepage	Y	76	String	Single	A URL such as http://example-example.com
WebVPN-Deny-Message	Y	116	String	Single	Valid string (up to 500 characters)
WebVPN-Download_Max-Size	Y	157	Integer	Single	0x7fffffff
WebVPN-File-Access-Enable	Y	94	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Browsing-Enable	Y	96	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Entry-Enable	Y	95	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Group-based-HTTP/HTTPS-Proxy-Exception-List	Y	78	String	Single	Comma-separated DNS/IP with an optional v (for example *.cisco.com, 192.168.1.*, wwwin
WebVPN-Hidden-Shares	Y	126	Integer	Single	0 = None 1 = Visible
WebVPN-Home-Page-Use-Smart-Tunnel	Y	228	Boolean	Single	Enabled if clientless home page is to be rende Smart Tunnel.
WebVPN-HTML-Filter	Y	69	Bitmap	Single	1 = Java ActiveX 2 = Scripts 4 = Image 8 =
WebVPN-HTTP-Compression	Y	120	Integer	Single	0 = Off 1 = Deflate Compression

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
WebVPN-HTTP-Proxy-IP-Address	Y	74	String	Single	Comma-separated DNS/IP:port, with http= or ht prefix (for example http=10.10.10.10:80, https=11.11.11.11:443)
WebVPN-Idle-Timeout-Alert-Interval	Y	148	Integer	Single	0-30. 0 = Disabled.
WebVPN-Keepalive-Ignore	Y	121	Integer	Single	0-900
WebVPN-Macro-Substitution	Y	223	String	Single	Unbounded.
WebVPN-Macro-Substitution	Y	224	String	Single	Unbounded.
WebVPN-Port-Forwarding-Enable	Y	97	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-Exchange-Proxy-Enable	Y	98	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-HTTP-Proxy	Y	99	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-List	Y	72	String	Single	Port forwarding list name
WebVPN-Port-Forwarding-Name	Y	79	String	Single	String name (example, “Corporate-Apps”). This text replaces the default string, “Application A on the clientless portal home page.
WebVPN-Post-Max-Size	Y	159	Integer	Single	0x7fffffff
WebVPN-Session-Timeout-Alert-Interval	Y	149	Integer	Single	0-30. 0 = Disabled.
WebVPN Smart-Card-Removal-Disconnect	Y	225	Boolean	Single	0 = Disabled 1 = Enabled
WebVPN-Smart-Tunnel	Y	136	String	Single	Name of a Smart Tunnel
WebVPN-Smart-Tunnel-Auto-Sign-On	Y	139	String	Single	Name of a Smart Tunnel auto sign-on list appen the domain name
WebVPN-Smart-Tunnel-Auto-Start	Y	138	Integer	Single	0 = Disabled 1 = Enabled 2 = Auto Start
WebVPN-Smart-Tunnel-Tunnel-Policy	Y	227	String	Single	One of “e networkname,” “i networkname,” or “a networkname is the name of a Smart Tunnel netw e indicates the tunnel excluded, i indicates the tu specified, and a indicates all tunnels.
WebVPN-SSL-VPN-Client-Enable	Y	103	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Keep- Installation	Y	105	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Required	Y	104	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSO-Server-Name	Y	114	String	Single	Valid string
WebVPN-Storage-Key	Y	162	String	Single	

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
WebVPN-Storage-Objects	Y	161	String	Single	
WebVPN-SVC-Keepalive-Frequency	Y	107	Integer	Single	15-600 seconds, 0=Off
WebVPN-SVC-Client-DPD-Frequency	Y	108	Integer	Single	5-3600 seconds, 0=Off
WebVPN-SVC-DTLS-Enable	Y	123	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SVC-DTLS-MTU	Y	125	Integer	Single	MTU value is from 256-1406 bytes.
WebVPN-SVC-Gateway-DPD-Frequency	Y	109	Integer	Single	5-3600 seconds, 0=Off
WebVPN-SVC-Rekey-Time	Y	110	Integer	Single	4-10080 minutes, 0=Off
WebVPN-SVC-Rekey-Method	Y	111	Integer	Single	0 (Off), 1 (SSL), 2 (New Tunnel)
WebVPN-SVC-Compression	Y	112	Integer	Single	0 (Off), 1 (Deflate Compression)
WebVPN-UNIX-Group-ID (GID)	Y	222	Integer	Single	Valid UNIX group IDs
WebVPN-UNIX-User-ID (UIDs)	Y	221	Integer	Single	Valid UNIX user IDs
WebVPN-Upload-Max-Size	Y	158	Integer	Single	0x7fffffff
WebVPN-URL-Entry-Enable	Y	93	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-URL-List	Y	71	String	Single	URL list name
WebVPN-User-Storage	Y	160	String	Single	
WebVPN-VDI	Y	163	String	Single	List of settings

Table 106: RADIUS Attributes Sent to Secure Firewall Threat Defense

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Address-Pools	217	String	Single	The name of a network object defined on the Firewall Threat Defense device that identifies a subnet, which will be used as the address pool for clients connecting to the remote access VPN. Define the network object on the **Objects** page.
Banner1	15	String	Single	The banner to display when the user logs in.
Banner2	36	String	Single	The second part of the banner to display when the user logs in. Banner2 is appended to Banner1.

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Downloadable ACLs	Cisco-AV-Pair	merge-dacl {before-avpair | after-avpair}		Supported via Cisco-AV-Pair configuration.
Filter ACLs	86, 87	String	Single	Filter ACLs are referred to by ACL name in the RADIUS server. It requires the ACL configuration to be already present on the Firewall Threat Defense device, so that it can be used during RADIUS authorization. 86=Access-List-Inbound 87=Access-List-Outbound
Group-Policy	25	String	Single	The group policy to use in the connection. You must create the group policy on the remote access VPN **Group Policy** page. You can use one of the following formats: • *group policy name* • OU=*group policy name* • OU=*group policy name*;
Simultaneous-Logins	2	Integer	Single	The number of separate simultaneous connections the user is allowed to establish, 0 - 2147483647.
VLAN	140	Integer	Single	The VLAN on which to confine the user's connection, 0 - 4094. You must also configure this VLAN on a subinterface on the Firewall Threat Defense device.

You must set the values of the IE-Proxy-Server-Method attribute returned from ISE to one of the following:

- IE_PROXY_METHOD_PACFILE: 8
- IE_PROXY_METHOD_PACFILE_AND_AUTODETECT: 11
- IE_PROXY_METHOD_PACFILE_AND_USE_SERVER: 12
- IE_PROXY_METHOD_PACFILE_AND_AUTODETECT_AND_USE_SERVER: 15

Firewall Threat Defense will deliver a proxy setting only if one of the above values is used for the IE-Proxy-Server-Method attribute.

Create or Update Aliases for a Connection Profile

Aliases contain alternate names or URLs for a specific connection profile. Remote access VPN administrators can enable or disable the Alias names and Alias URLs. VPN users can choose an Alias name when they connect to the Secure Firewall Threat Defense device. Aliases names for all connections configured on this device can be turned on or off for display. You can also configure the list of Alias URLs, which your endpoints

can select while initiating the remote access VPN connection. If users connect using the Alias URL, system will automatically log them using the connection profile that matches the Alias URL.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the policy that you want to modify.

Step 3 Click **Edit** on the connection profile for which you want to create or update aliases.

Step 4 Click **Aliases**.

Step 5 To add an Alias name, do the following:

a) Click **Add** under **Alias Names**.
b) Specify the **Alias Name**.
c) Select the **Enabled** check box in each window to enable the aliases.
d) Click **OK**.

Step 6 To add an Alias URL, do the following:

a) Click **Add** under **URL Alias**.
b) Select the **Alias URL** from the list or create a new URL object. For more information see Create a URL object, on page 3039.
c) Select the **Enabled** check box in each window to enable the aliases.
d) Click **OK**.

Step 7 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1832

Configure Access Interfaces for Remote Access VPN

The **Access Interface** table lists the interface groups and security zones that contain the device interfaces. These are configured for remote access SSL or IPsec IKEv2 VPN connections. The table displays the name of each interface group or security-zone, the interface trustpoints used by the interface, and whether Datagram Transport Layer Security (DTLS) is enabled.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Click the **Access Interface** tab.

Step 4 To add an access interface, click + and specify values for the following in the **Add Access Interface** dialog box:

a) **Access Interface**—Select the interface group or security zone to which the interface belongs.

The interface group or security zone must be a Routed type. Other interface types are not supported for remote access VPN connectivity.

b) Associate the **Protocol** object with the access interface by selecting the following options:

- **Enable IPSet-IKEv2**—Select this option to enable **IKEv2** settings.
- **Enable SSL**—Select this option to enable **SSL** settings.
  - Select **Enable Datagram Transport Layer Security**.

    When selected, it enables Datagram Transport Layer Security (DTLS) on the interface and allows the AnyConnect VPN module of Cisco Secure Client to establish an SSL VPN connection using two simultaneous tunnels—an SSL tunnel and a DTLS tunnel.

    Enabling DTLS avoids the latency and bandwidth problems associated with certain SSL connections and improves the performance of real-time applications that are sensitive to packet delays.

    To configure SSL settings, and TLS and DTLS versions, see Configure SSL settings, on page 1279.

    To configure SSL settings for the AnyConnect VPN module of Cisco Secure Client, see Group policy options for Secure Client, on page 3065.
  - Select the **Configure Interface Specific Identity Certificate** check box and select **Interface Identity Certificate** from the drop-down list.

    If you do not select the Interface Identity Certificate, the **Trustpoint** will be used by default.

    If you do not select the Interface Identity Certificate or Trustpoint, the **SSL Global Identity Certificate** will be used by default.

c) Click **OK** to save the changes.

Step 5 Select the following under **Access Settings**:

- **Allow Users to select connection profile while logging in**—If you have multiple connection profiles, check this check box to allow user to select the correct connection profile during login. You must select this option for **IPsec-IKEv2** VPNs.

Step 6 Use the following options to configure **SSL Settings**:

- **Web Access Port Number**—The port to use for VPN sessions. The default port is 443.
- **DTLS Port Number**—The UDP port to use for DTLS connections. The default port is 443.
- **SSL Global Identity Certificate**— The selected **SSL Global Identity Certificate** will be used for all the associated interfaces if the **Interface Specific Identity Certificate** is not provided.

Step 7 For **IPsec-IKEv2 Settings**, select the **IKEv2 Identity Certificate** from the list or add an identity certificate.

Step 8 Configure **Service Access Control**. Choose a service access object from the **Service Access Object** drop-down list or click + to create a new object.

You can use a service access object to control remote clients' access to VPN on Firewall Threat Defense devices with Version 7.7 or later. This object provides geolocation-based access control to clients before VPN authentication. By default, there is no access control for RA VPN, and remote clients can connect from any geolocation unless specified by a service access object. For more information, see Manage VPN Access of Remote Users Based on Geolocation, on page 1886and Configure a service access object, on page 2952.

Step 9 Under the **Access Control for VPN Traffic** section, select the following option if you want to bypass access control policy:

- **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** — Decrypted traffic is subjected to Access Control Policy inspection by default. Enabling the Bypass Access Control policy for decrypted traffic option bypasses the ACL inspection, but VPN Filter ACL and authorization ACL downloaded from AAA server are still applied to VPN traffic.

  Note
  If you select this option, you need not update the access control policy for remote access VPN as specified in Update the Access Control Policy on the Secure Firewall Threat Defense Device, on page 1825.

Step 10 Click **Save** to save the access interface changes.

Related Topics

Interfaces, on page 2980

Configure Advanced Options for Remote Access VPN

Cisco Secure Client Image

Secure Client Image

The Secure Client provides secure SSL or IPsec (IKEv2) connections to the Firewall Threat Defense device for remote users with full VPN profiling to corporate resources. Without a previously-installed client, remote users can enter the IP address of an interface configured to accept clientless VPN connections in their browser to download and install the Secure Client. The Firewall Threat Defense device downloads the client that matches the operating system of the remote computer. After downloading, the client installs and establishes a secure connection. In case of a previously installed client, when the user authenticates, the Firewall Threat Defense device, examines the version of the client, and upgrades the client if necessary.

The Remote Access VPN administrator associates any new or additional Secure Client images to the VPN policy. The administrator can unassociate the unsupported or end of life client packages that are no longer required.

The Cloud-Delivered Firewall Management Center determines the type of operating system by using the file package name. If the user renamed the file without indicating the operating system information, the valid operating system type must be selected from the list box.

Download the Secure Client image file by visiting Cisco Software Download Center.

Related Topics

Adding a Secure Client Image to the Cloud-Delivered Firewall Management Center, on page 1853

Adding a Secure Client Image to the Cloud-Delivered Firewall Management Center

You can upload the Secure Client image to the Cloud-Delivered Firewall Management Center by using the **Secure Client File** object. For more information, see File objects, on page 3075. For more information about the client image, see Cisco Secure Client Image, on page 1853.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**, choose and edit a listed remote access policy, then choose the **Advanced** tab.

Step 2 Click **Add** to add a Secure Client image.

Step 3 Click **Add** in the **Available Secure Client Images** portion of the **Secure Client Images** dialog.

Step 4 Enter the **Name** and **Description**(optional) for the available Secure Client Image.

Step 5 Click **Browse**, locate and select the client image that you want to upload.

Step 6 Click **Save** to upload the image to the Cloud-Delivered Firewall Management Center.

When you upload the client image to the Cloud-Delivered Firewall Management Center, the operating system information for the image appears automatically.

Step 7 To change the order of client images, Click **Show Re-order buttons** and move the client image up or down.

Related Topics

Cisco Secure Client Image, on page 1853

Update Secure Client Image for Remote Access VPN Clients

When new Secure Client updates are available in Cisco Software Download Center, you can download the packages manually and add them to the remote access VPN policy so that the new client packages are upgraded on the VPN client systems according to their operating systems.

Before you begin

Instructions in this section help you update new Secure Client images to remote access VPN clients connecting to Secure Firewall Threat Defense VPN gateway. Ensure that the following configurations are complete before updating your Secure Client images:

- Download the latest Secure Client image files from Cisco Software Download Center.
- On your Cloud-Delivered Firewall Management Center web interface, go to **Objects** > **Object Management** > **VPN** > **Secure Client File** and add the new Secure Client image files.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy that you want to update.

Step 3 From **Advanced** drop-down menu, choose **AnyConnect Client ImagesSecure Client Image** and then click **Add**.

Step 4 Select a client image file from **Available Secure Client Images** and click **Add**.

If the required client image is not listed, click **Add** to browse and upload an image.

Step 5 Click **OK**.

Step 6 Save the remote access VPN policy.
After the remote access VPN policy changes are deployed, the new Secure Client images are updated on the Secure Firewall Threat Defense device that is configured as the remote access VPN gateway. When a new VPN user connects to the VPN gateway, the user gets the new Secure Client image to download depending on the operating system of the client system. For existing VPN users, the Secure Client image gets updated in their next VPN session.

Add a Cisco Secure Client External Browser Package to the Cloud-Delivered Firewall Management Center

If you have the Secure Client external browser package image on your local disk, use this procedure to upload the same to the Cloud-Delivered Firewall Management Center. After you upload the external browser package, you can update the external browser package for your remote access VPN connections.

You can upload the external browser package file to the Cloud-Delivered Firewall Management Center by using the **Secure Client File** object. For more information, see File objects, on page 3075.

Points to Remember

- Only one external browser package can be added to the Firewall Threat Defense device.
- After the external browser package is added to the Cloud-Delivered Firewall Management Center, the browser is pushed to the Firewall Threat Defense only after the external browser is enabled in the remote access VPN configuration.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**, choose and edit a listed remote access policy, then choose the **Advanced** tab.

Step 2 Click **Add** in the **Secure Client External Browser Package** portion of the **Secure Client Images** page.

Step 3 Enter the **Name** and **Description** for the Secure Client package.

Step 4 Click **Browse** and locate the external browser package file to upload.

Step 5 Click **Save** to upload the image to the Cloud-Delivered Firewall Management Center.

Note
If you want to update the remote access VPN connection with an existing external browser package, select the file from the **Package File** drop-down.

Step 6 Save the remote access VPN policy.

Related Topics

Cisco Secure Client Image, on page 1853

Remote Access VPN Address Assignment Policy

The Firewall Threat Defense device can use an IPv4 or IPv6 policy for assigning IP addresses to Remote Access VPN clients. If you configure more than one address assignment method, the Firewall Threat Defense device tries each of the options until it finds an IP address.

IPv4 or IPv6 Policy

You can use the IPv4 or IPv6 policy to address an IP address to remote access VPN clients. You must try with the IPv4 policy to begin and later followed by IPv6 policy.

- **Use Authorization Server**—Retrieves the address from an external authorization server on a per-user basis. If you are using an authorization server that has IP address configured, we recommend using this method. Address assignment is supported by RADIUS-based authorization server only. It is not supported for AD/LDAP. This method is available for both IPv4 and IPv6 assignment policies.
- **Use DHCP**—Obtains IP addresses from a DHCP server configured in a connection profile. You can also define the range of IP addresses that the DHCP server can use by configuring DHCP network scope in the group policy. If you use DHCP, configure the server in the **Objects** > **Object Management** > **Network** pane. This method is available for IPv4 assignment policies.

  For more information about DHCP network scope configuration, see Group policy general options, on page 3063.
- **Use an internal address pool**—Internally configured address pools are the easiest method of address pool assignment to configure. If you use this method, create the IP address pools in the **Objects** > **Object Management** > **Address Pools** pane and select the same in the connection profile. This method is available for both IPv4 and IPv6 assignment policies.
- **Allow reuse an IP address so many minutes after it is released**—Delays the reuse of an IP address after its return to the address pool. Adding a delay helps to prevent problems firewalls can experience when an IP address is reassigned quickly. By default, the delay is set to zero. If you want to extend the delay, enter the number of minutes in the range of 0–480 to delay the IP address reassignment. This configurable element is available for IPv4 assignment policies.

Related Topics

Configure Connection Profile Settings, on page 1832
Remote Access VPN Authentication, on page 1814

Configure Certificate Maps

Certificate maps let you define rules matching a user certificate to a connection profile based on the contents of the certificate fields. Certificate maps provide certificate authentication on secure gateways.

The rules or the certificate maps are defined in Certificate map objects, on page 3058.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Choose **Advanced** > **Certificate Maps**.

Step 4 Select the following options from the **General Settings for Connection Profile Mapping** pane:

Selections are priority-based, matching continues down the list of options when the first selection does not match. Matching is complete when the rules are satisfied. If the rules are not satisfied, the default connection profile listed at the bottom of this page is used for the connection. Select any, or all of the following options to establish authentication and to determine which connection profile (tunnel group) must be mapped to the client.

- **Use Group URL if Group URL and Certificate Map match different Connection profiles**

- **Use the configured rules to match a certificate to a Connection Profile**—Enable this to use the rules defined in the Connection Profile Maps.

Note
Configuring a certificate mapping implies certificate-based authentication. The remote user will be prompted for a client certificate regardless of the configured authentication method.

Step 5 Under the **Certificate to Connection Profile Mapping** section, click **Add Mapping** to create certificate to connection profile mapping for this policy.

a) Choose or create a **Certificate Map Name** object.
b) Select the **Connection Profile** that want to use if the rules in the certificate map object are satisfied.
c) Click **OK** to create the mapping.

Step 6 Click **Save**.

Configuring Group Policies

A group policy is a set of attribute and value pairs, stored in a group policy object, that define the remote access VPN experience. For example, in the group policy object, you configure general attributes such as addresses, protocols, and connection settings.

The group policy applied to a user is determined when the VPN tunnel is established. The RADIUS authorization server assigns the group policy, or it is obtained from the current connection profile.

Note There is no group policy attribute inheritance on the Firewall Threat Defense. A group policy object is used, in its entirety, for a user. The group policy object identified by the AAA server upon login is used, or, if that is not specified, the default group policy configured for the VPN connection is used. The provided default group policy can be set to your default values, but will only be used if it is assigned to a connection profile and no other group policy has been identified for the user.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.
Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.
Step 3 Choose **Advanced** > **Group Policies** > **Add**.
Step 4 Select group policies from the **Available Group Policy** list and click **Add**. You can select one or more group policies to associate with this remote access VPN policy.
Step 5 Click **OK** to complete the group policy selection.
Step 6 Save your changes.

Related Topics

Configure a group policy object, on page 3062

Configuring LDAP Attribute Mapping

An LDAP attribute name maps LDAP user or group attribute name to a Cisco-understandable name. The attribute map equates attributes that exist in the Active Directory (AD) or LDAP server with Cisco attribute names. You can map any standard LDAP attribute to a well-known vendor specific attribute (VSA). You can map one or more LDAP attributes to one or more Cisco LDAP attributes. When the AD or LDAP server returns authentication to the Firewall Threat Defense device during remote access VPN connection establishment, the Firewall Threat Defense device can use the information to adjust how the Secure Client completes the connection.

When you want to provide VPN users with different access permissions or VPN content, you can configure different VPN policies on the VPN server and assign these policy-sets to each user based on their credentials. You can achieve this in Firewall Threat Defense by configuring LDAP authorization, with LDAP attribute maps. In order to use LDAP to assign a group policy to a user, you must configure a map that maps an LDAP attribute.

An LDAP attribute map consists of three components:

- **Realm**—Specifies the name for the LDAP attribute map; the name is generated based on the selected realm.
- **Attribute Name Map**—Maps the LDAP user or group attribute name to Cisco-understandable name.
- **Attribute Value Map**—Maps value in the LDAP user or group attribute to the value of a Cisco attribute for the selected name mapping.

The group policies that are used in an LDAP attribute map get added to the list of group policies in the remote access VPN configuration. Removing a group policy from the remote access VPN configuration also removes the associated LDAP attribute mapping.

In versions 6.4 to 6.6, you can configure LDAP attribute maps only using FlexConfig. For more information, see Configure AnyConnect Modules and Profiles Using FlexConfig.

In versions 7.0 and later, you can use the following procedure:

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Click **Advanced** > **LDAP Attribute Mapping**.

Step 4 Click **Add**.

Step 5 On the Configure LDAP Attribute Map page, select a **Realm** to configure the attribute map.

Step 6 Click **Add**.

You can configure multiple attribute maps. Each attribute map requires that you configure a name map and value maps.

Note

Ensure that you follow these guidelines while creating an LDAP attribute map:

- Configure at least one mapping for an LDAP attribute; multiple mappings with the same LDAP attribute name is not allowed.

- Configure a minimum of one name map to create an LDAP attribute map.
- You can remove any LDAP attribute map if the attribute map is not associated with any connection profile in the remote access VPN configuration.
- Use the correct spelling and capitalization in the LDAP attribute map for ***both*** the Cisco and LDAP attribute names and values.

a) Specify the **LDAP Attribute Name** and then select the required **Cisco Attribute Name** from the list.
b) Click **Add Value Map** and Specify the **LDAP Attribute Value** and **Cisco Attribute Value**.

Repeat this step to add more value maps.

Step 7 Click **OK** to complete LDAP attribute map configuration.

Step 8 Click **Save** to save the changes to the LDAP attribute mapping.

Example

For a detailed example, see Configure RA VPN with LDAP Authentication and Authorization for FTD.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1835
Understanding Policy Enforcement of Permissions and Attributes, on page 1816

Configuring VPN Load Balancing

About VPN Load Balancing

VPN load balancing in Firewall Threat Defense allows you group two or more devices logically and distribute remote access VPN sessions among the devices equally. VPN load balancing shares Secure Client VPN sessions among the devices in a load balancing group.

VPN load balancing is based on simple distribution of traffic without taking into account throughput or other factors. A VPN load-balancing group consists of two or more Firewall Threat Defense devices. One device acts as the director, and the other devices are member devices. Devices in a group do not need to be of the exact same type, or have identical software versions or configurations. Any Firewall Threat Defense device that supports remote access VPN can participate in a load balancing group. Firewall Threat Defense supports VPN load balancing with Secure Client SAML authentication.

All active devices in a VPN load-balancing group carry session loads. VPN load balancing directs traffic to the least-loaded device in the group, distributing the load among all devices. It makes efficient use of system resources and provides increased performance and high availability.

Components of VPN Load Balancing

Following are the components of VPN load balancing:

- **Load-balancing group**—A virtual group of two or more Firewall Threat Defense devices to share the VPN sessions.

A VPN load-balancing group can consist of Firewall Threat Defense devices of the same release or of mixed releases; but the device must support remote access VPN configuration.

See Configure Group Settings for VPN Load Balancing, on page 1860 and Configure Additional Settings for Load Balancing, on page 1861.

- **Director**—One device from the group acts a director. It distributes the load among other members in the group and participate is serving the VPN sessions.

  The director monitors all devices in the group, keeps track of how loaded each device is, and distributes the session load accordingly. The role of director is not tied to a physical device; it can shift among devices. For example, if the current director fails, one of the member devices in the group takes over that role and immediately becomes the new director.

- **Members**—Devices other than the director in a group are called members. They participate in load balancing and share the remote access VPN connections.

  Configure Settings for Participating Devices, on page 1862.

Prerequisites for VPN Load Balancing

- **Certificates**—Firewall Threat Defense's certificate must contain the IP addresses or FQDN of the director and members to which the connection is redirected. Or else, the certificate will be deemed untrusted. The certificate must use Subject Alternate Name (SAN) or wildcard certificate
- **Group URL**—Add the group URL for VPN load-balancing group IP address to the connection profiles. Specify a group URL to eliminate the need for the user to select a group at login.
- **IP Address Pool**—Choose unique IP address pool for member devices, and override the IP pool in Cloud-Delivered Firewall Management Center for each of the member devices.
- Devices that are behind Network Address Translation (NAT) can also be part of a load balancing group.

Guidelines and Limitations for VPN Load Balancing

- VPN load balancing is disabled by default. You must explicitly enable VPN load balancing.
- Only the Firewall Threat Defense devices that are co-located can be added to a load-balancing group.
- A load-balancing group must have a minimum of two Firewall Threat Defense devices.
- Devices in Firewall Threat Defense high availability can participate in a load-balancing group.
- Devices that are behind Network Address Translation (NAT) can also be part of a load balancing group.
- When a member or a director device goes down, remote access VPN connections that are served by that device will be dropped. You must initiate the VPN connection again.
- Identity certificate on each device must have Subject Alternate Name (SAN) or wildcard.

Configure Group Settings for VPN Load Balancing

You can enable VPN load balancing and configure group settings that are applicable to all the members of the load-balancing group. When you create the group, you can configure participation settings for load balancing.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy that you want to update.

Step 3 Click **Advanced** > **Load Balancing**.

Step 4 Click the **Enable Load balancing between member devices** toggle button to enable load balancing. The **Edit Group Configuration** page opens. Group parameters apply to all devices under the load-balancing group.

Step 5 Specify the **Group IPv4 Address** and **Group IPv6 Address** as applicable.

The IP address that you specify here is for the entire load-balancing group and the director opens this IP address for incoming VPN connections.

Step 6 Select the **Communication Interface** for the load-balancing group. Click **Add** to add an interface group or security zone.

Communication interface is a private interface through which the director and members share information about their load.

Step 7 Enter the **UDP Port** for communication between the director and members in a group. The default port is 9023.

Step 8 Enable the **IPsec Encryption** toggle button to activate IPsec encryption for the communication between the director and members.

Enabling the encryption establishes an IPsec tunnel between the director and members using a pre-shared key.

When you upgrade or downgrade Firewall Threat Defense devices with the **IPsec Encryption** option enabled, ensure there is no configuration mismatch between the Cloud-Delivered Firewall Management Center and the Firewall Threat Defense to prevent deployment failures.

Step 9 Enter **Encryption Key** for IPsec encryption and confirm the encryption key.

Step 10 Click **OK**.

Configure Additional Settings for Load Balancing

The additional settings for VPN load balancing include FQDN and IKEv2 redirection.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy that you want to update.

Step 3 Click **Advanced** > **Load Balancing**.

Step 4 Turn on the **Enable Load balancing between member devices** toggle button to enable load balancing if not done already.

Step 5 Click **Settings**.

Step 6 Turn on the **Send FQDN to peer devices instead of IP** toggle button to enable redirection using a fully qualified domain name.

By default, Firewall Threat Defense sends only IP addresses in VPN load balancing redirection to a client.

Step 7 Select one of the **IKEv2 Redirect** phases:

- **Redirect during SA authentication**
- **Redirect during SA initialization**

Step 8 Click **OK**.

Step 9 Save your changes.

Configure Settings for Participating Devices

The device participation settings determine how the devices share load in VPN load balancing. Configure a participating device by enabling VPN load balancing on the device and defining device-specific properties. These values vary from device to device. You can provide a priority number for the devices participating in load balancing. A higher priority number gives the device a better chance to become the director over other devices. But you cannot select a device to be the director of the group.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** next to the remote access VPN policy that you want to modify.

Step 3 Click **Advanced** > **Load Balancing**.

Step 4 Turn on the **Enable Load balancing between member devices** toggle button to enable load balancing if you have not enabled already.

Step 5 Configure **Device Participation** settings:

The **Device Participation** section lists all the target devices of the selected remote access VPN configuration. You can configure these devices to share the load of the incoming VPN sessions.

a) Turn on the **Load Balancing** toggle button to enable load balancing for a device and then click **Edit**.

b) Enter the device **Priority**.

By default, the device priority is set to 5. You can choose a number from 1 trough 10.

c) Specify the **IPv4 NAT** or **IPv6 NAT** address for VPN interface IP address if the device is behind NAT.

d) Click **OK**.

Step 6 Click **Save** to save the remote access VPN policy settings.

Configuring IPsec Settings for Remote Access VPNs

The IPsec settings are applicable only if you selected IPsec as the VPN protocol while configuring your remote access VPN policy. If not, you can enable IKEv2 using the Edit Access Interface dialog box. See Configure Access Interfaces for Remote Access VPN, on page 1851 for more information.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 From the list of available VPN policies, select the policy for which you want to modify the settings.

Step 3 Click **Advanced**.

The list of IPsec settings appears in a navigation pane on the left of the screen.

Step 4 Use the navigation pane to edit the following IPsec options:

a) **Crypto Maps**—The Crypto Maps page lists the interface groups on which IKEv2 protocol is enabled. Crypto Maps are auto generated for the interfaces on which IKEv2 protocol is enabled. To edit a Crypto Map, see Configure Remote Access VPN Crypto Maps, on page 1863. You can add or remove interface groups to the selected VPN policy in **Access Interface**. See Configure Access Interfaces for Remote Access VPN, on page 1851 for more information.

b) **IKE Policy**—The IKE Policy page lists all the IKE policy objects applicable for the selected VPN policy when Secure Client endpoints connect using the IPsec protocol. See IKE Policies in Remote Access VPNs, on page 1865 for more information. To add a new IKE policy, see Configure an IKEv2 policy object, on page 3073. Firewall Threat Defense supports only Secure Client IKEv2 clients. Third-party standard IKEv2 clients are not supported.

c) **IPsec/IKEv2 Parameters**—The IPsec/IKEv2 Parameters page enables you to modify the IKEv2 session settings, IKEv2 Security Association settings, IPsec settings, and NAT Transparency settings. See Configure Remote Access VPN IPsec/IKEv2 Parameters, on page 1866 for more information.

Step 5 Click **Save**.

Configure Remote Access VPN Crypto Maps

Crypto maps are automatically generated for the interfaces on which IPsec-IKEv2 protocol has been enabled. You can add or remove interface groups to the selected VPN policy in **Access Interface**. See Configure Access Interfaces for Remote Access VPN, on page 1851 for more information.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 From the list of available VPN policies, select the policy for which you want to modify the settings.

Step 3 Click the **Advanced** > **Crypto Maps**, and select a row in the table and click **Edit** to edit the Crypto map options.

Step 4 Select **IKEv2 IPsec Proposals** and select the transform sets to specify which authentication and encryption algorithms will be used to secure the traffic in the tunnel.

Step 5 Select **Enable Reverse Route Injection** to enable static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint.

Step 6 Select **Enable Client Services** and specify the port number.

The Client Services Server provides HTTPS (SSL) access to allow the Secure Client Downloader to receive software upgrades, profiles, localization and customization files, CSD, SCEP, and other file downloads required

by the client. If you select this option, specify the client services port number. If you do not enable the Client Services Server, users will not be able to download any of these files that the Secure Client might need.

Note
You can use the same port that you use for SSL VPN running on the same device. Even if you have an SSL VPN configured, you must select this option to enable file downloads over SSL for IPsec-IKEv2 clients.

Step 7 Select **Enable Perfect Forward Secrecy** and select the **Modulus group**.

Use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the **Modulus Group** list.

Modulus group is the Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select the modulus group that you want to allow in the remote access VPN configuration:

- 1—Diffie-Hellman Group 1 (768-bit modulus).
- 2—Diffie-Hellman Group 2 (1024-bit modulus).
- 5—Diffie-Hellman Group 5 (1536-bit modulus, considered good protection for 128-bit keys, but group 14 is better). If you are using AES encryption, use this group (or higher).
- 14—Diffie-Hellman Group 14 (2048-bit modulus, considered good protection for 128-bit keys).
- 19—Diffie-Hellman Group 19 (256-bit elliptical curve field size).
- 20—Diffie-Hellman Group 20 (384-bit elliptical curve field size).
- 21—Diffie-Hellman Group 21 (521-bit elliptical curve field size).
- 24—Diffie-Hellman Group 24 (2048-bit modulus and 256-bit prime order subgroup).

Step 8 Specify the **Lifetime Duration (seconds)**.

The lifetime of the security association (SA), in seconds. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. Generally, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes.

You can specify a value from 120 to 2147483647 seconds. The default is 28800 seconds.

Step 9 Specify the **Lifetime Size (kbytes)**.

The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires.

You can specify a value from 10 to 2147483647 kbytes. The default is 4,608,000 kilobytes. No specification allows infinite data.

Step 10 Select the following **ESPv3 Settings**:

- **Validate incoming ICMP error messages**—Choose whether to validate ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.

- **Enable 'Do Not Fragment' Policy**—Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header, and select one of the following from the **Policy** list:
  - Copy—Maintains the DF bit.
  - Clear—Ignores the DF bit.
  - Set—Sets and uses the DF bit.
- Select **Enable Traffic Flow Confidentiality (TFC) Packets**— Enable dummy TFC packets that mask the traffic profile which traverses the tunnel. Use the **Burst**, **Payload Size**, and **Timeout** parameters to generate random length packets at random intervals across the specified SA.

  Note
  Enabling traffic flow confidentiality (TFC) packets prevents the VPN tunnel from being idle. Thus the VPN idle timeout configured in the group policy does not work as expected when you enable the TFC packets. See Advanced group policy option fields and values, on page 3068.

  - Burst—Specify a value from 1 to 16 bytes.
  - Payload Size—Specify a value from 64 to 1024 bytes.
  - Timeout—Specify a value from 10 to 60 seconds.

Step 11 Click **OK**.

Related Topics

Interfaces, on page 2980

IKE Policies in Remote Access VPNs

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs). The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection. An IKE proposal is a set of algorithms that two peers use to secure the negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters are used to protect subsequent IKE negotiations.



Note Firewall Threat Defense supports only IKEv2 for remote access VPNs.

Unlike IKEv1, in an IKEv2 proposal, you can select multiple algorithms and modulus groups in one policy. Since peers choose during the Phase 1 negotiation, this makes it possible to create a single IKE proposal, but consider creating multiple, different proposals to give higher priority to your most desired options. For IKEv2, the policy object does not specify authentication, other policies must define the authentication requirements.

An IKE policy is required when you configure a remote access IPsec VPN.

Configuring Remote Access VPN IKE Policies

The IKE Policy table specifies all the IKE policy objects applicable for the selected VPN configuration when Secure Client endpoints connect using the IPsec protocol. For more information, see IKE Policies in Remote Access VPNs, on page 1865.

Note Firewall Threat Defense supports only IKEv2 for remote access VPNs.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 From the list of available VPN policies, select the policy for which you want to modify the settings.

Step 3 Click **Advanced** > **IKE Policy**.

Step 4 Click **Add** to select from the available IKEv2 policies, or add a new IKEv2 policy and specify the following:

- **Name**—Name of the IKEv2 policy.
- **Description**—Optional description of the IKEv2 policy
- **Priority**—The priority value determines the order of the IKE policy compared by the two negotiating peers when attempting to find a common security association (SA).
- **Lifetime**— Lifetime of the security association (SA), in seconds
- **Integrity**—The Integrity Algorithms portion of the Hash Algorithm used in the IKEv2 policy.
- **Encryption**—The Encryption Algorithm used to establish the Phase 1 SA for protecting Phase 2 negotiations.
- **PRF Hash**—The pseudorandom function (PRF) portion of the Hash Algorithm used in the IKE policy. In IKEv2, you can specify different algorithms for these elements.
- **DH Group**—The Diffie-Hellman group used for encryption.

Step 5 Click **Save**.

Configure Remote Access VPN IPsec/IKEv2 Parameters

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 From the list of available VPN policies, select the policy for which you want to modify the settings.

Step 3 Click **Advanced** > **IPsec**> **IPsec/IKEv2 Parameters**.

Step 4 Select the following for **IKEv2 Session Settings**:

- **Identity Sent to Peers**—Choose the identity that the peers will use to identify themselves during IKE negotiations:
  - **Auto**—Determines the IKE negotiation by connection type: IP address for preshared key, or Cert DN for certificate authentication (not supported).
  - **IP address**—Uses the IP addresses of the hosts exchanging ISAKMP identity information.
  - **Hostname**—Uses the fully qualified domain name (FQDN) of the hosts exchanging ISAKMP identity information. This name comprises the hostname and the domain name.
- **Enable Notification on Tunnel Disconnect**—Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA does not match the traffic selectors for that SA. Sending this notification is disabled by default.
- **Do not allow device reboot until all sessions are terminated**—Check to enable waiting for all active sessions to voluntarily terminate before the system reboots. This is disabled by default.

Step 5 Select the following for **IKEv2 Security Association (SA) Settings**:

- **Cookie Challenge**—Whether to send cookie challenges to peer devices in response to SA initiated packets, which can help thwart denial of service (DoS) attacks. The default is to use cookie challenges when 50% of the available SAs are in negotiation. Select one of these options:
  - **Custom**—Specify **Threshold to Challenge Incoming Cookies**, the percentage of the total allowed SAs that are in-negotiation. This triggers cookie challenges for any future SA negotiations. The range is zero to 100%. The default is 50%.
  - Always— Select to send cookie challenges to peer devices always.
  - Never— Select to never send cookie challenges to peer devices.
- **Number of SAs Allowed in Negotiation**—Limits the maximum number of SAs that can be in negotiation at any time. If used with Cookie Challenge, configure the cookie challenge threshold lower than this limit for an effective cross-check. The default is 100 %.
- **Maximum number of SAs Allowed**—Limits the number of allowed IKEv2 connections.

Step 6 Select the following for **IPsec Settings**:

- **Enable Fragmentation Before Encryption**—This option lets traffic travel across NAT devices that do not support IP fragmentation. It does not impede the operation of NAT devices that do support IP fragmentation.
- **Path Maximum Transmission Unit Aging**—Check to enable PMTU (Path Maximum Transmission Unit) Aging, the interval to Reset PMTU of an SA (Security Association).
- **Value Reset Interval**—Enter the number of minutes at which the PMTU value of an SA (Security Association) is reset to its original value. Valid range is 10 to 30 minutes, default is unlimited.

Step 7 Select the following for **NAT Settings**:

- **Keepalive Messages Traversal**—Select whether to enable NAT keepalive message traversal. NAT traversal keepalive is used for the transmission of keepalive messages when there is a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow. If you select this option, configure the interval, in seconds, between the keepalive signals sent

between the spoke and the middle device to indicate that the session is active. The value can be from 10 to 3600 seconds. The default is 20 seconds.

- **Interval**—Sets the NAT keepalive interval, from 10 to 3600 seconds. The default is 20 seconds.

Step 8 Click **Save**.

Customize Cisco Secure Client

You can configure the Secure Client customizations using the management center and deploy them to the VPN headend. The threat defense device distributes these customizations to the endpoint when a user connects from the Secure Client. Administrators can customize the following elements at the endpoint:

Table 107: Secure Client Customizations Available in Management Center

Customization	Description	More Info
GUI text and messages	Customize or localize Secure Client GUI text and informational/error messages.	Customize and Localize Secure Client GUI Text and Messages, on page 1872
Icons and images	Customize the logo, images, or icons of the Secure Client GUI.	Customize Secure Client Icons and Images, on page 1873
Scripts	Deploy scripts on the endpoint devices when a client establishes or disconnects a VPN session.	Deploy Scripts on Endpoint Devices Using Secure Client, on page 1874
Binaries	Deploy custom applications using Secure Client APIs.	Deploy Custom Applications Using Cisco Secure Client APIs, on page 1876
Custom Installer Transforms	Customize the Secure Client installer.	Customize the Secure Client Installer, on page 1877
Localized Installer Transforms	Localize the Secure Client installer.	Localize the Client Installer, on page 1877

Overview of the customization file

The customization file is an XML-based configuration object in the Threat Defense device that controls the visual appearance and layout of Secure Client to remote VPN users. It defines:

- Branding—logos and color schemes
- Page structure—logon page and portal layout
- Content elements—titles, banners, and button labels
- Messaging—error messages and other user-facing text

By default, a Threat Defense device has two customization objects:

- Template: A blank XML template file that you can update to add the required customizations. You cannot delete this file.
- DfltCustomization: An XML file with a set of URLs and basic customization to bring up the Secure Client. This is the default customization file.

  To view these customization files, run the **show import webvpn customization** command on the Threat Defense device:

```
FTD(config)# show import webvpn customization
Template
DfltCustomization
```

Customize the customization file

To update the **Template** XML file:

1. Export the file to disk0. You can rename the file, if required.

```
FTD(config)# export webvpn customization Template disk0:ABCTemplate
```

2. Copy the XML file to a remote PC.
3. Edit the XML file with the required customizations and copy it back to disk0.
4. Import the file to the Threat Defense device.

```
FTD(config)# import webvpn customization ABCTemplate disk0:ABCTemplate
```

5. Verify the customization files in the device.

```
FTD(config)# show import webvpn customization
Template
DfltCustomization
ABCTemplate
```



Note You can also update the **DfltCustomization** file using the same procedure as above. You can also rename this file.

Upload the customization file to Management Center

1. Choose **Objects** > **Object Management** > **VPN** > **Secure Client File** and click **Add Secure Client File**.
2. In the **Add Secure Client File** dialog box, configure these parameters:
   a. In the **Name** file, enter the name of the customization file.
   b. From the **File Type** drop-down list, choose **Secure Client VPN Profile**.
   c. Click **Browse**, locate and select the customization file that you want to upload.
   d. Click **Save**.
3. Choose **Devices** > **VPN** > **Remote Access**.
4. Click the edit icon next to the remote access VPN policy.

5. Click the **Advanced** tab.
6. In the left pane, click **Group policies**.
7. Click **+** to add a group policy.
8. In the **Add Group Policy** dialog box, configure these parameters:
   a. Enter a **Name** and optionally a **Description** for this policy.
   b. Click the **Secure Client** tab, and from the **Client Profile** drop-down list, choose the customization file.
   c. Click **Save**.
9. Choose this group policy as the selected policy in the **Group Policy** dialog box, click **OK**, and click **Deploy**.

Guidelines and Limitations for Secure Client Customizations

Guidelines

- Ensure that you import the xml file to the object defined as `DfltCustomization`. When you import the customization object, the Cloud-Delivered Firewall Management Center checks the XML code for validity.

  For example:

  ```
  hostname# import webvpn customization DfltCustomization
  disk0:/csm/defaultcustomizationwebvpn.xml
  ```

- Add the Windows Cisco Secure Client headend deployment package to the remote access VPN policy before deploying customization commands. This package must be included when supporting Windows, MacOS, or Linux clients.

General Limitations

- If there are any customizations earlier than 7.4 configured directly on the threat defense using the CLI, the management center removes these customizations during deployment.
- No clustering support.

Table 108: Limitations for Secure Client Customizations

Customization	Limitations

GUI Text and Message Customizations	• You must restart Secure Client before you use this customization.
	• No support for right-to-left languages.
	• Some strings are truncated in the GUI due to hardcoded field lengths.
	• Some messages are hardcoded in the client. For example:
	  • Status messages (during an update)
	  • Untrusted server messages
	  • Deferred update messages
	• Localization is version-specific.
	If an Administrator created a translation table based on a template of an old Secure Client version, the new messages do not appear to the remote users. The Administrator must merge the latest template with the translation table so that the table has the new messages. You can use third-party tools like Get text to perform the merge.

Icons and Images Customizations	• You must restart Secure Client before you use this customization. • The customization object names must match the Secure Client GUI filenames. File names are different for each operating system and are case-sensitive. For more information about the filenames, extensions, and sizes for each OS, see the Cisco Secure Client Administrator Guide. • Images do not appear properly if the image size is not correct. The management center or the threat defense device does not verify the image size. • MacOS does not support customization of Secure Client images and icons.
Deploy Customized Scripts	• Secure Client runs only one OnConnect and one OnDisconnect script; however, these scripts may launch other scripts. • Scripts can only execute functions that the user has the right to invoke. • You cannot launch the OnConnect script from the Start Before Logon (SBL) GUI.
Deploy Custom Applications Using Cisco Secure Client APIs (Binaries)	• After you use this customization, if you deploy an updated version of the Secure Client on the management center, the client will download the update and replace your custom UI.

Customization	Limitations
Customize the Client Installer	• This customization is available only for Windows.

Customize and Localize Secure Client GUI Text and Messages

You can customize the Secure Client GUI text and the informational/error messages. You also can customize the GUI text and all messages to appear in a preferred language.

Customize GUI Text and Messages

To customize the GUI text or messages, you can edit the messages in the message file. You can update the messages to include more details in the error messages. For example:

- Modify any label in the login dialog box such as **Password** to **Domain Password**.
- Add your support contact details in the error messages.

Localize GUI Text and Messages

Threat defense devices use translation tables to translate the labels and user messages displayed in the Secure Client. When a user connects to the remote access VPN, the Secure Client identifies the locale set on the endpoint and downloads the translation file. Threat defense devices support 128 locales. By default, Secure Client is installed in English.

- For Cisco Secure Client 5.0, the default localization files for various languages are part of the application.
- For AnyConnect Clients 4.x, you can download the localization files for a few languages from Cisco.com and upload them to the management center.

How to Customize Secure Client GUI Text and Messages

Before you begin

Configure one or more remote access VPN policies.

Procedure

Step 1 Get the base template or translation file either from the Secure Client package or from Cisco.com. For example, AnyConnect.po.

Step 2 Use a text editor to add the translations or customize the labels or messages.

Step 3 Update the **msgstr** string corresponding to each **msgid**.

Step 4 Save the file.

Step 5 Create a new Secure Client Customization object.

a) Choose **Objects** > **Object Management** > **VPN** > **Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization.
d) From the **Customization Type** drop-down list, choose **GUI Text and Messages**.
e) From the **Language** drop-down list, choose a language for which you are adding the translation.

f) Click **Browse** and select the translation file. The supported file extensions are .po, .mo, and .txt.

Step 6 Add the customization to the remote access VPN policy.

a) Choose **Devices** > **VPN** > **Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced > Secure Client Customizations > Localization**.
d) Click **+** to select the translation file.
e) Click **Add**.
f) Click **OK**.

Step 7 Click **Save**.

Step 8 Choose **Deploy > Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

Verify the Secure Client customizations. For more information, see Verify Secure Client Customizations, on page 1878.

Customize Secure Client Icons and Images

You can customize the logo, images, and icons of the Secure Client GUI using the management center.

To customize the logo, images, and icons of the Secure Client GUI, you must create a Secure Client customization object in the management center. The name of this customization object must match the Secure Client GUI filename. File names are different for each operating system and are case-sensitive. For more information about the filenames, extensions, and sizes for each OS, see the Cisco Secure Client Administrator Guide.

For example, if you want to replace the corporate logo for Windows clients, you must create a customization object with the name as `company_logo` and add this customization object to the remote access VPN policy. If you use a different name for the customization object, the Secure Client installer does not change the component. However, if you deploy your own executable to customize the GUI, the customization object can have any name.

Location of all images files:

- Windows: **%PROGRAMFILES%\Cisco\Cisco Secure Client\res**
- Linux: **/opt/cisco/secureclient/resources** or **/opt/cisco/anyconnect/resources**
- macOS: Not supported



Note When the remote access VPN connection is established, the logo or images are downloaded to the client directory. Ensure that you restart the client for the updated logo or images to appear in your client.

How to Customize Secure Client Images and Icons

Before you begin

Configure one or more remote access VPN policies.

Procedure

Step 1 Create the icons or images using the correct extensions, and sizes.

Images do not appear properly if the image size is not correct. The management center or the threat defense device does not verify the image size.

For more information about the filenames, extensions, and sizes, see the Cisco Secure Client Administrator Guide.

Step 2 Create a new Secure Client Customization object.

a) Choose **Objects** > **Object Management** > **VPN** > **Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization object.

Ensure that the customization object name matches the Secure Client GUI filenames. For more information about the filenames, see the Cisco Secure Client Administrator Guide.

d) From the **Customization Type** drop-down list, choose **Icon and Images**.
e) From the **Platform** drop-down list, choose a platform.
f) Click **Browse** and select the file. The supported file extensions are .png, .ico, and .jpeg.
g) Repeat the steps from 2a - f to add multiple icons and images.

Step 3 Add the customization object to the remote access VPN policy.

a) Choose **Devices** > **VPN** > **Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced > Secure Client Customizations > Icons and Images**.
d) Click **+** to select the file.
e) Click **Add**.
f) Click **OK**.

Step 4 Click **Save**.

Step 5 Choose **Deploy > Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

Verify the Secure Client customizations. For more information, see Verify Secure Client Customizations, on page 1878.

Deploy Scripts on Endpoint Devices Using Secure Client

Secure Client lets you download and run scripts when the following events occur:

- Establishment of a new client VPN session with the threat defense. The script triggered by this event is an OnConnect script because it requires this filename prefix. Reconnection of the VPN session does not trigger this script.
- Disconnection of a client VPN session with the threat defense. The script triggered by this event is an OnDisconnect script because it requires this filename prefix.

These scripts run asynchronously and do not delay the connection establishment or disconnection. They can be of any extension and must be executable on the endpoint. Secure Client identifies the OnConnect and OnDisconnect scripts by the filename. It looks for a file whose name begins with OnConnect or OnDisconnect regardless of the file extension.

Some examples of this feature are:

- Refresh the group policy upon VPN connection.
- Mount a network drive upon a VPN connection.
- Unmount a network drive upon a VPN disconnection.

To enable scripts, check the **Enable Scripting** option in the VPN profile. By default, the client does not launch scripts. The client does not require the script to be written in a specific language. It requires an application that can run the script to be installed on the client computer. For the client to launch the script, the script must run from the command line.

Secure Client can only launch scripts after the user logs in and establishes a VPN session. You cannot launch the OnConnect script from the Start Before Logon (SBL) GUI. You must check the **Enable Post SBL On Connect Script** option in the VPN profile to trigger the scripts once the user logs in. Secure Client is a 32-bit application. When the scripts run on a 64-bit Windows version, they use the 32-bit version of cmd.exe.

How to Add Customized Scripts for Secure Client

Before you begin

1. Configure a remote access VPN.
2. Enable scripting in the VPN profile.
3. Add the VPN profile to the remote access VPN group policy.

Procedure

Step 1 Create OnConnect and OnDisconnect scripts for a platform.

Step 2 Create a new Secure Client Customization object.

a) Choose **Objects** > **Object Management** > **VPN** > **Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization.
d) From the **Customization Type** drop-down list, choose **Scripts**.
e) From the **Platform** drop-down list, choose a platform.
f) Select one of the following:

- **On Connect**: To select an OnConnect script.

- **On Disconnect**: To select an OnDisconnect script.

g) Click **Browse** and select the script you want to execute on the endpoint.

Step 3 Add the customization to the remote access VPN policy.

a) Choose **Devices** > **VPN** > **Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced > Secure Client Customizations > Localization**.
d) Click **+** to select the translation file.
e) Click **Add**.
f) Click **OK**.

Step 4 Click **Save**.

Step 5 Choose **Deploy > Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

Verify the Secure Client customizations. For more information, see Verify Secure Client Customizations, on page 1878.

Deploy Custom Applications Using Cisco Secure Client APIs

For Windows, Linux, or MacOs machines, you can create and deploy a custom client that uses the Secure Client APIs. You can use this client binary file to replace the Secure Client GUI or CLI binary files.

Your executable can call any resource files such as logo images that you import to the management center. When you deploy your own executable, you can use any filenames for your resource files.

The following table lists the filenames of the client executable files for the different operating systems:

Table 109: Filenames of the Secure Client Executable Files

Client OS	Client GUI File	Client CLI File
Windows	vpnui.exe	vpncli.exe
Linux	vpnui	vpn
MacOS	Not supported by the management center deployment. However, you can deploy an executable for the macOS that replaces the client GUI using other means, such as Altiris Agent.	vpn

How to Deploy Custom Applications Using Cisco Secure Client API

Before you begin

Configure one or more remote access VPN policies.

Procedure

Step 1 Create the custom application using the Cisco Secure Client APIs.

Step 2 Create a new Secure Client Customization object.

a) Choose **Objects** > **Object Management** > **VPN** > **Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization.
d) From the **Customization Type** drop-down list, choose **Binary**.
e) From the **Platform** drop-down list, choose a platform.
f) Click **Browse** and select the custom application.

Step 3 Add the customization to the remote access VPN policy.

a) Choose **Devices** > **VPN** > **Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced > Secure Client Customizations > Binaries**.
d) Click **+** to select the translation file.
e) Click **Add**.
f) Click **OK**.

Step 4 Click **Save**.

Step 5 Choose **Deploy > Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Customize the Secure Client Installer

You can customize the Secure Client GUI by creating a custom transform that deploys with the client installer program. You can import the transform to the management center and deploy it to the threat defense device. The threat defense deploys this transform to the endpoint using the client installer program.

Note This customization is available only for Windows.

Localize the Client Installer

You can translate the messages displayed by the Cisco Secure Client installer. The management center uses a transform to translate the messages displayed by the installer. The transform alters the installation but leaves the original security-signed MSI intact. These transforms only translate the installer screens and do not translate the client GUI screens.

Every release of Cisco Secure Client includes a localized transform that administrators can upload to the management center whenever they upload Cisco Secure Client packages with new software. If you use our localization transform, ensure that you update them with the latest release from Cisco.com whenever you upload a new Cisco Secure Client package.

How to Customize or Localize the Client Installer

Before you begin

Configure one or more remote access VPN policies.

Procedure

Step 1 Create a customized or localized transform.

Step 2 Create a new Secure Client Customization object.

a) Choose **Objects** > **Object Management** > **VPN** > **Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization.
d) From the **Customization Type** drop-down list, choose **Customized Installer Transform** or **Localized Installer Transform**.
e) From the **Platform** drop-down list, choose a platform.
f) Click **Browse** and select the transform.

Step 3 Add the customization to the remote access VPN policy.

a) Choose **Devices** > **VPN** > **Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced > Secure Client Customizations > Custom Installer Transforms** or **Localized Installer Transforms**.
d) Click **+** to select the transform.
e) Click **Add**.
f) Click **OK**.

Step 4 Click **Save**.

Step 5 Choose **Deploy > Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Verify Secure Client Customizations

Validate the Deployment

After the deployment is complete, validate the deployment in the management center. Click the **Transcript Details** ( )) icon to verify the commands generated for the customizations.

Examples

1. The example below shows the transcript details of a GUI text and messages customization: Secure Client en_us localization with customized prompts.

```
import webvpn translation-table AnyConnect language en-us disk0:/AnyConnect_en-us.po
```

2. The example below shows the transcript details of customized icons and images: customization of the Secure Client logo to ABC logo.

```
import webvpn AnyConnect-customization type resource platform win name company_logo.png
 disk0:/company_logo.png
```

3. The example below shows the transcript details of the customized OnConnect and OnDisconnect scripts. The On connect script mounts the network drive and the On disconnect script unmounts the network drive.

```
import webvpn AnyConnect-customization type binary platform win name
  scripts_OnConnect_mount.bat disk0:/mount.bat
import webvpn AnyConnect-customization type binary platform win name
  scripts_OnDisconnect_unmount.bat disk0:/unmount.bat
```

Verify Customizations Using Threat Defense Commands

Use the following commands on the threat defense to verify the customizations:

- **show import webvpn translation-table detailed**: Shows the available translation tables.

```
HQ-FTD# show import webvpn translation-table detailed
 Translation Tables' Templates:
   AnyConnect           ia4DaAXNSvl5pZboQRGJcs9KMXY=
   customization
 Translation Tables:
   fr                   customization        BWWodsOt1PbvDvYOp8hLb3W7a64=
   ja                   customization        lNvUk1+qTLNZyNrBcApMQPHnm1M=
   ru                   customization        UqyKyUAcjR+xTGUtdiIFnoIiW5U=
```

- **show import webvpn AnyConnect-customization detailed**: Shows the details of the Secure Client customizations.

```
HQ-FTD# show import webvpn AnyConnect-customization detailed
OEM resources for AnyConnect client:
  linux-64/binary/scripts_OnConnect_conn.sh           w6+n7z80D/8AR+ul2f7DvTmcDTw=
  linux-64/binary/scripts_OnDisconnect_discon.sh      jx5LJC2XBEmEkGeww59CAkszvnI=
  linux-64/resource/company-logo.png                  GsfBDroqGSQEEwuBDS/3DJNVv88=
  win/binary/scripts_OnConnect_mount.bat              dzjfsLYYft/XMlPlzskKl+Wv1bw=
  win/binary/scripts_OnDisconnect_unmount.bat         k6xlKhF1l2IRyJu08+sdYXgKNgM=
  win/resource/company_logo.png                       cmEvxwqvtaS+Pz/6sb9n3NZudS4=
```

Verify Customizations in Secure Client

- In Secure Client, click the **Message History** tab to verify that the customizations are downloaded.

  Use the DART tool to view the client-side diagnostics.

Table 110: Verifying the Customizations in Secure Client

Customization	Verification
GUI text and message customizations	• Verify if Secure Client has the language localization or customized files in the following locations: • Windows: %ProgramData%\Cisco\Cisco AnyConnect Secure Mobility Client\l10n\<language-code>\LC_MESSAGES (AnyConnect versions 4.9 and earlier) • Windows: %ProgramData%\Cisco\Cisco Secure Client\l10n\<language-code>\LC_MESSAGES (Secure Client versions 5.0 and later) • For Mac OS and Linux: /opt/cisco/anyconnect/l10n/<LANGUAGE-CODE>/LC_MESSAGES /opt/cisco/secureclient/l10n/<LANGUAGE-CODE>/LC_MESSAGES • Verify the contents of the localized or customized file and confirm if they have the customized or localized strings. For example: AnyConnect.mo
Image and icon customizations	• Verify if Secure Client has the customized files in the following locations: • Windows: %PROGRAMFILES%\ Cisco\Cisco AnyConnect Secure Mobility Client\res\ (AnyConnect versions 4.10 and earlier) • Windows: %PROGRAMFILES%\ Cisco\Cisco Secure Client\UI\res (Cisco Secure Client 5.0 and above) • Linux: /opt/cisco/anyconnect/resources or /opt/cisco/secureclient/resources • Verify the content of the customized icon or image file.
Customized OnConnect and OnDisconnect scripts	• Verify if the customized scripts are in the following locations: • Windows: %ProgramData%\Cisco\Cisco Secure Client\Script (Cisco Secure Client 5.0 and above) • Windows: %ProgramData%\Cisco\ Cisco AnyConnect Secure Mobility Client\Script (AnyConnect versions 4.9 and earlier) • Mac OS and Linux: /opt/cisco/anyconnect/script or /opt/cisco/secureclient/vpn/script • Verify the script.

Configure Secure Client Management VPN Tunnel

A management VPN tunnel provides connectivity to the corporate network whenever a client system is powered up, without the VPN users having to connect to the VPN. This helps organizations keep their endpoints up-to-date with software patches and updates. Management tunnel disconnects when the user-initiated VPN tunnel is established.

This section provides information about configuring Secure Client management VPN tunnel on Firewall Threat Defense. Configuring the Secure Client management tunnel on Firewall Threat Defense using the Cloud-Delivered Firewall Management Center web interface requires the following settings:

- A **Connection profile** with certificate-based authentication and a group URL.
- **Secure Client management VPN profile file**, configured a server with group URL and backup servers if required.
- A **Group policy** with the management VPN profile, split tunneling with explicitly included networks, client bypass protocol, and no banner.

For detailed instructions on configuring the Secure Client Management VPN tunnel, see Configuring Secure Client Management VPN Tunnel on Firewall Threat Defense, on page 1882.

Requirements and Prerequisites for Secure Client Management VPN Tunnel

Software and Configuration Requirements

Ensure that you have the following before you configure the Secure Client Management tunnel on using the Firewall Threat Defense using the Cloud-Delivered Firewall Management Center web interface:

- Ensure that you are using Firewall Threat Defense and Cloud-Delivered Firewall Management Center versions 6.7.0 or above.
- Download the Secure ClientSecure Client VPN Webdeploy package 4.7 or above and upload it to Firewall Threat Defense remote access VPN.
- Ensure that the certificate authentication is configured in the connection profile.
- Ensure that no banner is configured in the group policy.
- Check the split tunneling configuration in the management tunnel-group policy.

Certificate Requirements

- Firewall Threat Defense must have a valid identity certificate for remote access VPN and the root certificate from the local certifying authority (CA) must be present on the Firewall Threat Defense.
- Endpoints connecting to the management VPN tunnel must have a valid identity certificate.
- CA certificate for Firewall Threat Defense's identity certificate must be installed on the endpoints and the CA certificate for the endpoints must be installed on the Firewall Threat Defense.
- The identity certificate issued by the same local CA must be present in the Machine store.

  Certificate Store (For Windows) and/or in System Keychain (For macOS).

Limitations of Secure Client Management VPN Tunnel

- Secure Client Management VPN Tunnel supports only certificate authentication, it does not support AAA-based authentication.
- Public or private proxy settings are not supported.
- Secure Client upgrade and AnyConnect module download are not supported when the management VPN tunnel is connected.

Configuring Secure Client Management VPN Tunnel on Firewall Threat Defense

Procedure

Step 1 **Create a remote access VPN policy configuration using the wizard**:

For information about configuring a remote access VPN, see Configuring a New Remote Access VPN Connection, on page 1821.

Step 2 **Configure connection profile settings for management VPN tunnel**:

Note
It is advisable to create a new connection profile to be used only for Secure Client management VPN tunnel.

a) Edit the remote access VPN policy you have created.
b) Select and edit the connection profile that will be used for management VPN tunnel.
c) From **AAA** drop-down menu, choose **Authentication Method** and click **Client Certificate Only**. Configure the authorization and accounting settings as required.
d) Click the **Aliases** tab of the connection profile.
e) Click **Add (+)** under URL Aliases and **URL Alias** for the connection profile.
f) Click **Enabled** to enable the URL.
g) Click **OK** and then click **Save** to save the connection profile settings.

For more information about connection profile settings, see Configure Connection Profile Settings, on page 1832.

Step 3 **Create a management tunnel profile using the Secure Client profile editor**:

a) Download the Secure Client **VPN Management Tunnel Standalone Profile Editor** from Cisco Software Download Center if you have not done already.
b) Create a management tunnel profile with the required settings for your VPN users and save the file.
c) Configure a server in the Server List with the group URL you have configured in the connection profile.

For information about creating a management profile using the Profile Editor, see the *Cisco Secure Client (including AnyConnect) Administrator Guide*.

Step 4 **Create a management tunnel object**:

a) On your Cloud-Delivered Firewall Management Center web interface, navigate to **Objects** > **Object Management** > **VPN** > **Secure Client File**
b) Click **Add Secure Client File**.
c) Specify the **Name** for the Secure Client file.

d) Click **Browse** and select the management tunnel profile file you have saved.
e) Click the **File Type** drop-down and select **Secure Client Management VPN Profile**.
f) Click **Save**.

Note
You an also create the management tunnel object when you create or update Secure Client settings for a group policy. See Group policy options for Secure Client, on page 3065.

Step 5 **Associate a management profile with a group policy and configure group policy settings**:

You must add the management VPN profile to the group policy associated with the connection profile used for the management tunnel VPN connection. When the user connects, the management VPN profile is downloaded along with the user VPN profile already mapped to the group policy, enabling the management VPN tunnel feature.

Caution
No Banner: Check and ensure that no banner is configured in the group policy settings. You can check the banner settings under **Group Policy** > **General Settings** > **Banner**.

a) Edit the connect profile you have created for management VPN tunnel.
b) Click **Edit Group Policy** > **Secure Client** > **Management Profile Edit Group Policy** and choose **AnyConnect Secure Client**, then click **Management Profile**.
c) Click the **Management VPN Profile** drop-down and select the management profile file object you have created.

Note
You can also click + and add a new Secure Client Management VPN Profile object.

d) Click **Save**.

Step 6 **Configure split tunneling in group policy**:

a) Click **Edit Group Policy** and then choose **General**, then click **Split Tunneling**.
b) From the IPv4 or IPv6 split tunneling drop-down, select **Tunnel networks specified below**.
c) Select the Split Tunnel Network List Type: **Standard Access List** or **Extended Access List**, and then select the required access list to allow the traffic over the management VPN tunnel.
d) Click **Save** to save the split tunnel settings.

Secure Client Custom Attribute

Secure Client Management VPN tunnel requires split include tunneling configuration by default. If you are configuring Secure Client custom attribute in the group policy to deploy the management VPN tunnel with split tunneling to tunnel all, you can do so using FlexConfig because Cloud-Delivered Firewall Management Center 6.7 web interface does not support Secure Client custom attribute .

The following is an example command for Secure Client custom attribute:

```
webvpn
 anyconnect-custom-attr ManagementTunnelAllAllowed description ManagementTunnelAllAllowed
anyconnect-custom-data ManagementTunnelAllAllowed true true
group-policy MGMT_Tunnel attributes
 anyconnect-custom ManagementTunnelAllAllowed value true
```

Step 7 **Deploy, verify, and monitor the remote access VPN policy**:

a) Deploy the management VPN tunnel configuration to Firewall Threat Defense.

Note

Client systems must connect to the Firewall Threat Defense remote access VPN once to download the management tunnel VPN profile to the client machines.

b) You can verify the Secure Client management VPN tunnel at **AnyConnect Secure Mobility ClientSecure Mobility Client** and choose **VPN**, then click **Statistics**.

You can also check the management VPN session details on the Firewall Threat Defense command prompt using the **show vpn-sessiondb anyconnect** command.

c) On your Cloud-Delivered Firewall Management Center web interface, click **Analysis** to view the management tunnel session information.

Related Topics

Configure Connection Profile Settings, on page 1832
Firewall Threat Defense group policy objects, on page 3061

Multiple Certificate Authentication

Multiple certificate based authentication gives the ability to have the Firewall Threat Defense validate the machine or device certificate, to ensure the device is a corporate-issued device, in addition to authenticating the user's identity certificate to allow VPN access using the Secure Client during SSL or IKEv2 EAP phase.

The multiple certificates option allows certificate authentication of both the machine and user via certificates. Without this option, you could only do certificate authentication of either machine or the user, but not both.

Guidelines and Limitations of Multiple Certificate Authentication

Note When you configure multiple certificate authentication, ensure that you set the value of **AutomaticCertSelection** to true in the Cisco Secure Client Profile settings.

- Multiple certificate authentication currently limits the number of certificates to two.
- Secure Client must indicate support for multiple certificate authentication. If that is not the case then the gateway uses one of the legacy authentication methods or fails the connection. Secure Client version 4.4.04030 or later supports Multi-Certificate based authentication.
- Secure Client supports only RSA-based certificates.
- Only SHA256, SHA384, and SHA512 based certificate are supported during the Secure Client aggregate authentication.
- Certificate authentication cannot be combined with SAML authentication.

Configuring Multiple Certificate Authentication

Before you begin

Before you configure multiple certificate authentication, ensure that you have configured the certificate enrollment object that is used to obtain the identity certificate for each Firewall Threat Defense device. For more information, see Certificate map objects, on page 3058.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select the remote access VPN policy and click **Edit**.

Note
If you have not configured a remote access VPN, click **Add** to create a new remote access VPN policy.

Step 3 Select and **Edit** a connection profile to configure multiple certificate authentication.

Step 4 Click **AAA** settings and select **Authentication Method** > **Client Certificate Only** or **Client Certificate & AAA**.

Note
Select the **Authentication Server** if you have selected the Client Certificate & AAA authentication method

Step 5 Select the **Enable multiple certificate authentication** checkbox.

Step 6 Choose one of the certificates to **Map username from client certificate**:

- **First Certificate**— Select this option to map the username from the machine certificate sent from the VPN client.
- **Second Certificate**— Select this option to map the username from the user certificate sent from the client.

The username sent from the client is used as the VPN session username when certificate only authentication is enabled. When AAA and certificate authentication is enabled, VPN session username will be based on prefill option.

Note
If you select the **Map specific field** option, which includes the username from the client certificate, the **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively.

If you select the **Use entire DN (Distinguished Name) as username** option, the system automatically retrieves the user identity. A distinguished name (DN) is a unique identification, made up of individual fields that can be used as the identifier when matching users to a connection profile DN rules are used for enhanced certificate authentication.

If you have selected the Client Certificate & AAA authentication, select the **Prefill username from certificate on user login window** option to prefill the secondary username from the client certificate when the user connects via AnyConnect VPN module of Cisco Secure Client.

- **Hide username in login window**: The secondary username is pre-filled from the client certificate, but hidden to the user so that the user does not modify the pre-filled username.

Step 7 Configure the required AAA settings and connection profile settings for the remote access VPN.

Step 8 Save the connection profile and remote access VPN configuration and deploy it on your Firewall Threat Defense device.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1835

Manage VPN Access of Remote Users Based on Geolocation

You can manage the remote access VPN connections of users based on their geolocations. By configuring rules such that they allow or deny VPN connections from specific countries or regions, you can meet compliance requirements and enhance security. Connections that don't meet these location-based criteria are blocked before authentication and the details are logged in the Cloud-Delivered Firewall Management Center (**Devices** > **Troubleshoot** > **Troubleshooting Logs**).

This feature supports all versions of Secure Client.

Prerequisites

Firewall Threat Defense device must be Version 7.7.0 or later.

Workflow to Manage VPN Access of Remote Users Based on Geolocation

Step	Task	More Information
1	Define a policy to allow geolocation-based access control for your remote clients.	Configure a service access object, on page 2952
2	Update the remote access VPN configuration with the policy.	Configure Access Interfaces for Remote Access VPN, on page 1851
3	Deploy the configurations on the devices.	—
4	After clients connect to the Firewall Threat Defense device: **1.** Verify the active remote access VPN sessions in the **Remote Access VPN** dashboard. **2.** Verify the logs for the denied remote access VPN sessions at **Devices** > **Troubleshoot** > **Troubleshooting Logs**.	Monitor and Troubleshoot Service Access Policies, on page 1887

Guidelines and Limitations for Managing Remote Access VPN Users Based on Geolocation

Guidelines

- In a service access object, if you use a geolocation object (country, continent, or geolocation object), use it only in one rule.
- Configure the service access rules in the correct order because you cannot reorder these rules.

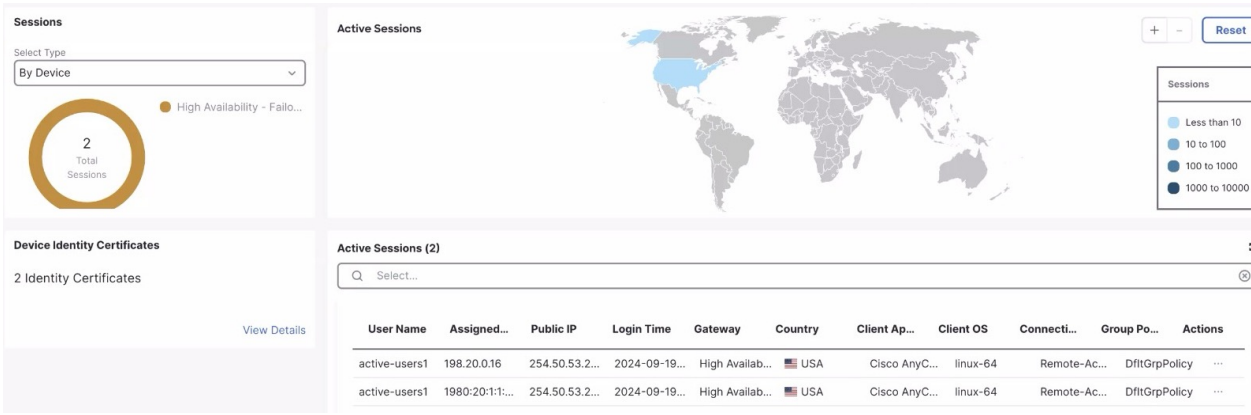
Limitations

- Clustering is not supported.
- Geolocation-based unclassified IP addresses are not categorized according to their geographic origin. For such unclassified IP addresses, the Cloud-Delivered Firewall Management Center enforces the default service access policy action.
- Connections from IETF RFC1918 addresses are allowed for geolocation-based remote access VPN regardless of service-access policy settings.

Monitor and Troubleshoot Service Access Policies

Monitor Active Remote Access VPN Sessions in Remote Access VPN Dashboard

Choose **Overview** > **Dashboards** > **Remote Access VPN**.

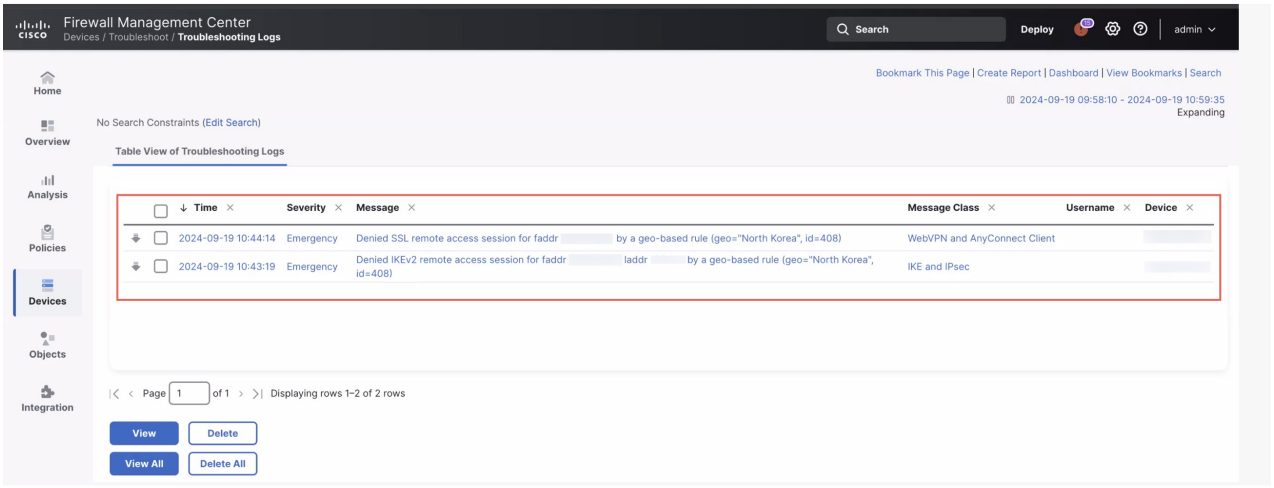


Monitor Denied Remote Access VPN Sessions

Monitor the denied remote access VPN sessions at **Devices** > **Troubleshoot** > **Troubleshooting Logs**. To view the denied remote access VPN sessions, you must configure the syslog settings in the Firewall Threat Defense:

1. Choose **Devices** > **Platform Settings** and create or edit a threat defense policy.
2. In the left pane, click **Syslog**.
3. Click the **Logging Setup** tab.
4. Check the **Enable Logging** check box.
5. Click the **VPN Logs** radio button.
6. From the **Logging Level** drop-down list, choose **6 - informational**.
7. Click **Save**.

Note You cannot view the denied remote access VPN sessions if the **All Logs** option is configured with a **Logging Level** between 0 and 2.



Verify Service Access Policies

From the Firewall Threat Defense device CLI, run the following commands:

- **show running-config service-access**: Displays the user-defined service access policies.

```
firepower#show running-config service-access
   service-access deny geolocation OBJGRP_Asia1
   service-access permit interface outside ra-ssl-client geolocation OBJGRP_India
   service-access deny ra-ikev2 geolocation any
```

- **show service-access**: Displays details of the user-defined service access policies.

```
firepower# show service-access
1 outside          : ra-ikev2 ra-ssl-client (permit) hits = 8288
  Last hit time    : 10:58:10.038 IST Tue Jul 16 2024
  object-group     : FMC_INTERNAL_XXY
2 any              : ra-ikev2 ra-ssl-client (deny) hits = 123
  Last hit time    : 11:23:12.032 IST Tue Jul 17 2024
  object-group     : any


firepower# show service-access detail
1 outside          : ra-ikev2 ra-ssl-client (permit) hits = 8288
  Last hit time    : 10:58:10.038 IST Tue Jul 16 2024
  object-group     : FMC_INTERNAL_XXY
  geolocation      : Egypt(818) Jordan(400)
                     Iran (Islamic Republic of)(364)
                     Saudi Arabia(682)
```

- **show geodb**: Displays details of the geolocation table.

  show geodb { **ipv4** | **ipv6** | **counters** | **context** } [**location** *country_name* | **lookup** *ip_address*] [**detail**]

  - **show geodb** { **ipv4** | **ipv6** } : Displays the total number of IPv4 or IPv6 address mappings.

```
firepower# show geodb ipv4
 Geolocation Table - IPv4
 Total number of mappings available: 532507
 Last geolocation data read time:  17:02:13.000 IST Thu Jul 18 2024
```

```
Running geolocation update version: 2024-02-15-019
```

- **show geodb** { **ipv4** | **ipv6** } **location** *country_name* **detail**: Displays the details of the IPv4 or IPv6 address mappings.

```
firepower# show geodb ipv4 location Antarctica detail
Geolocation Table - IPv4
id=0x00007fff82c284e0, geo_id=10, hits=0
        range_lower=77.70.176.176, range_upper=77.70.176.183
id=0x00007fff82cca360, geo_id:10, hits=0
     range_lower=79.110.169.69, range_upper=79.110.169.69
Total number of mappings available: 28
```

- **show geodb counters**: Displays the details of active, permitted, and denied sessions.

```
firepower# show geodb counters
current   - ongoing sessions
permitted - cumulative permitted sessions
denied    - cumulative denied sessions
Location                current          permitted      denied
Egypt                     0                  0             5
India                     45                 1345         45
```

- **show geodb** { **ipv4** | **ipv6lookup** *ip_address*: Displays the geolocation of a specific IPv4 or IPv6 address.

```
firepower# show geodb ipv4 lookup 223.223.128.24
Geolocation of 223.223.128.24 is "India" (356) with id=0x000015114d0aa330
Matching network range: 223.223.128.0 - 223.223.159.255
```

Troubleshoot Service Access Policies

- **Syslogs**

  Enable remote access VPN service access syslogs:

  1. Choose **Devices** > **Platform Settings**.
  2. Create or edit a platform settings policy.
  3. In the left pane, click **Syslog**.
  4. Click the **Logging Setup** tab and check the **Enable Logging** check box.
  5. Click the **Syslog Settings** tab and enable the syslogs for service access syslog 751031 and 716166.

- **Commands**
  - Use the **show running-config service-access**, and **show service-access** commands to view details of user-defined service access policies.
  - Use the **show geodb** command to view details of the geolocation table.
  - Use the **debug geolocation <debug-level>** command to capture debug logs related to geolocation. The debug levels can be 1 (Error), 2 (Warning), 3 and 4 (Info), 5 (Debug), or 255 (Debug all).

- Use the **clear geodb counters** command to clear the geolocation table counters such as the hit counts of the service access policies. However, you cannot clear the actual permitted and denied counters for locations using this command; you can clear these counters only after you reboot the device.

Customizing Remote Access VPN AAA Settings

This section provides information about customizing your AAA preferences for remote access VPNs. For more information, see Configure AAA Settings for Remote Access VPN, on page 1835.

Authenticate VPN Users via Client Certificates

You can configure remote access VPN authentication using client certificate when you create a new remote access VPN policy using the wizard or by editing the policy later.

Before you begin

Configure the certificate enrollment object that is used to obtain the identity certificate for each Firewall Threat Defense device that acts as a VPN gateway.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access policy and click **Edit**; or click **Add** to create a new remote access VPN policy.

Step 3 For a new remote access VPN policy, configure the authentication while selecting connection profile settings. For an existing configuration, select the connection profile that includes the client profile, and click **Edit**.

Step 4 Click **AAA** > **Authentication Method** > **Client Certificate Only**.

With this authentication method, the user is authenticated using a client certificate. You must configure the client certificate on VPN client endpoints. By default, the user name is derived from client certificate fields CN and OU respectively. If the user name is specified in other fields in the client certificate, use 'Primary' and 'Secondary' field to map appropriate fields.

If you select **Map specific field** option, which includes the username from the client certificate. The **Primary** and **Secondary** fields display the following default values, respectively: **CN (Common Name)** and **OU (Organisational Unit)**. If you select the **Use entire DN as username** option, the system automatically retrieves the user identity. A distinguished name (DN) is a unique identification, made up of individual fields, that can be used as the identifier when matching users to a connection profile. DN rules are used for enhanced certificate authentication.

- Primary and Secondary fields pertaining to the **Map specific field** option contain these common values:
  - C (Country)
  - CN (Common Name)
  - DNQ (DN Qualifier
  - EA (Email Address)

- GENQ (Generational Qualifier)
- GN (Given Name)
- I (Initial)
- L (Locality)
- N (Name)
- O (Organisation)
- OU (Organisational Unit)
- SER (Serial Number)
- SN (Surname)
- SP (State Province)
- T (Title)
- UID (User ID)
- UPN (User Principal Name)

- Whichever authentication method you choose, select or deselect **Allow connection only if user exists in authorization database**.

For more information, see Configure AAA Settings for Remote Access VPN, on page 1835.

Step 5 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1832
Add certificate enrollment objects, on page 2999

Configure VPN User Authentication via Client Certificate and AAA Server

When you configure remote access VPN authentication to use both client certificate and authentication server, VPN client authentication is done using both the client certificate validation and AAA server.

Before you begin

- Configure the certificate enrollment object that you use to obtain the identity certificate for each Firewall Threat Defense device that acts as a VPN gateway.
- Configure the RADIUS server group object and any AD or LDAP realms to use in the remote access VPN policy configuration.
- Ensure that the AAA Server is reachable from the Secure Firewall Threat Defense device for the remote access VPN configuration to work.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy for which you want to update the authentication or click **Add** to create new one.

Step 3 If you choose to create new remote access VPN policy, configure the authentication while selecting connection profile settings. For an existing configuration, select the connection profile that includes the client profile, and click **Edit**.

Step 4 Go to **AAA** and from the **Authentication Method** drop-down, choose **Client Certificate & AAA**.

- When you select the **Authentication Method** as:

  Client Certificate & AAA—Both types of authentication are done.

  - **AAA**—If you select the **Authentication Server** as **RADIUS**, by default, the Authorization Server has the same value. Select the **Accounting Server** from the drop-down list. Whenever you select **AD** and **LDAP** from the Authentication Server drop-down list, you must manually select the **Authorization Server** and **Accounting Server** respectively.
  - **Client Certificate**—Authenticates the user with client certificate. You must configure client certificate on the VPN client endpoints. By default, the username is derived from client certificate fields CN & OU respectively. If you use any other field in the client profile to specify the username, use **Primary Field** and **Secondary Field** to map appropriate fields.

    If you select **Map specific field** option, which includes the username from the client certificate. The **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively. If you select the **Use entire DN as username** option, the system automatically retrieves the user identity. A distinguished name (DN) is a unique identification, made up of individual fields that can be used as the identifier when matching users to a connection profile. DN rules are used for enhanced certificate authentication.

    Primary and Secondary fields pertaining to the **Map specific field** option contains these common values:

    - C (Country)
    - CN (Common Name)
    - DNQ (DN Qualifier
    - EA (Email Address)
    - GENQ (Generational Qualifier)
    - GN (Given Name)
    - I (Initial)
    - L (Locality)
    - N (Name)
    - O (Organisation)

- OU (Organisational Unit)
- SER (Serial Number)
- SN (Surname)
- SP (State Province)
- T (Title)
- UID (User ID)
- UPN (User Principal Name)

- Whichever authentication method you choose, select or deselect **Allow connection only if user exists in authorization database**.

For more information, see Configure AAA Settings for Remote Access VPN, on page 1835.

Step 5 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1832
Add certificate enrollment objects, on page 2999

Manage Password Changes over VPN Sessions

Password management allows remote access VPN policy administrator to configure the notification settings for the remote access VPN users on their password expiry. Password management is available in AAA settings with authentication methods AAA Only and Client Certificate & AAA. For more information, see Configure AAA Settings for Remote Access VPN, on page 1835.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy that you want to update.

Step 3 Click **Edit** on the connection profile that includes AAA settings.

Step 4 Choose **AAA** > **Advanced Settings** > **Password Management**.

Step 5 Check the **Enable Password Management** check-box and select one of the following:

- Notify User - days ahead of password expiry and specify the number of days in the box.
- Notify user on the day of password expiration.

Step 6 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1832

Send Accounting Records to the RADIUS Server

Accounting records in remote access VPN help the VPN administrator track the services that users access and the amount of network resources that they consume. Accounting information includes when user session start and stop, username, the number of bytes that pass through the device for each session, the service used, and the duration of each session.

You can use accounting alone or together with authentication and authorization. When you activate AAA accounting, the network access server reports the user activity to the configured accounting server. You can configure a RADIUS server as the accounting server so that the Cloud-Delivered Firewall Management Center sends all the user activity information to the RADIUS server.

Note You can use the same RADIUS server or separate RADIUS servers for authentication, authorization, and accounting in remote access VPN AAA settings.

Before you begin

- Configure a RADIUS group object with RADIUS servers to receive authentication requests or accounting records. For more information, see RADIUS server group configuration options, on page 2944.
- Ensure that the RADIUS server is reachable from the Firewall Threat Defense device. Configure routing on your Cloud-Delivered Firewall Management Center at **Devices** > **Device Management**, and click **Edit** ( ). Then, click the **Routing** tab to ensure connectivity to the RADIUS server.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access policy for which you want to configure RADIUS server, or create new remote access VPN policy.

Step 3 Click **Edit** on the connection profile that includes AAA settings and choose **AAA**.

Step 4 Select the RADIUS server from the **Accounting Server** drop-down.

Step 5 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1832
Configure AAA Settings for Remote Access VPN, on page 1835

Delegating Group Policy Selection to Authorization Server

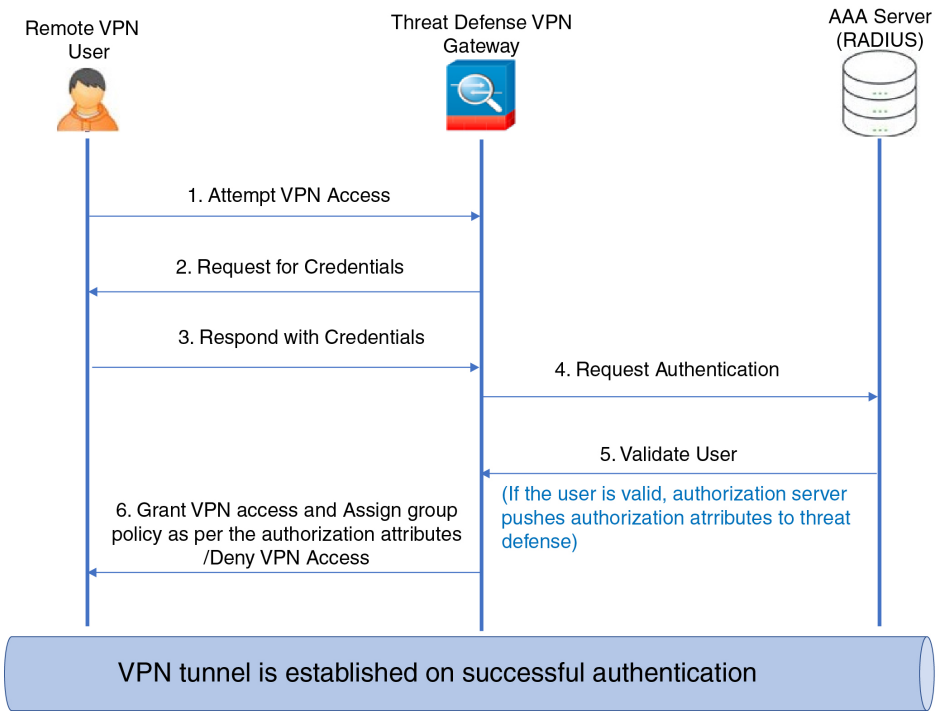
The group policy applied to a user is determined when the VPN tunnel is being established. You can select a group policy for a connection profile while creating a remote access VPN policy using the wizard or update the connection policy for connection profiles later. However, you can configure the AAA (RADIUS) server to assign the group policy or it is obtained from the current connection profile. If the Firewall Threat Defense

device receives attributes from the external AAA server that conflicts with those configured on the connection profile, then attributes from the AAA server always take the precedence.

You can configure ISE or the RADIUS Server to set the Authorization Profile for a user or user-group by sending IETF RADIUS Attribute 25 and map to the corresponding group policy name. You can configure specific group policy to a user or user group to push a Downloadable ACL, set a banner, Restrict VLAN, and configure the advanced option of applying an SGT to the session. These attributes are applied to all users that are part of that group when the VPN connection is established.

For more information, see the Configure Standard Authorization Policies section of Cisco Identity Services Engine Administrator Guide and RADIUS Server Attributes for Secure Firewall Threat Defense, on page 1841.

Figure 374: Remote Access VPN Group Policy Selection by AAA Server



Related Topics

Configure a group policy object, on page 3062
Configure Connection Profile Settings, on page 1832

Override the Selection of Group Policy or Other Attributes by the Authorization Server

When a remote access VPN user connects to the VPN, the group policy and other attributes configured in the connection profile are assigned to the user. However, the remote access VPN system administrator can delegate the selection of group policy and other attributes to the authorization server by configuring ISE or the RADIUS Server to set the Authorization Profile for a user or user-group. Once users are authenticated, these specific authorization attributes are pushed to the Firewall Threat Defense device.

Before you begin

Ensure that you configure a remote access VPN policy with RADIUS as the authentication server.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access policy and click **Edit**.

Step 3 Select RADIUS or ISE as the authorization server if not configured already.

Step 4 Select **Advanced** > **Group Policies** and add the required group policy. For detailed information about a group policy object, see Configure a group policy object, on page 3062.

You can map only one group policy to a connection profile; but you can create multiple group policies in a remote access VPN policy. These group policies can be referenced in ISE or the RADIUS server and configured to override the group policy configured in the connection profile by assigning the authorization attributes in the authorization server.

Step 5 Deploy the configuration on the target Firewall Threat Defense device.

Step 6 On the authorization server, create an Authorization Profile with RADIUS attributes for IP address and downloadable ACLs.

When the group policy is configured in the authorization server selected for remote access VPN, the group policy overrides the group policy configured in the connection profile for the remote access VPN user after the user is authenticated.

Related Topics

Configure a group policy object, on page 3062

Deny VPN Access to a User Group

When you do not want an authenticated user or user group to be able to use VPN, you can configure a group policy to deny VPN access. You can configure a group policy in a remote access VPN policy and reference it in the ISE or RADIUS server configuration for authorization.

Before you begin

Ensure that you have configured remote access VPN using the Remote Access Policy wizard and configured authentication settings for the remote access VPN policy.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access policy and click **Edit**.

Step 3 Select **AdvancedGroup Policies**.

Step 4 Select a group policy and click **Edit** or add a new group policy.

Step 5 Select **AdvancedSession Settings** and set **Simultaneous Login Per User** to 0 (zero).
This stops the user or user group from connecting to the VPN even once.

Step 6 Click **Save** to save the group policy and then save the remote access VPN configuration.

Step 7 Configure ISE or the RADIUS server to set the Authorization Profile for that user/user-group to send IETF RADIUS Attribute 25 and map to the corresponding group policy name.

Step 8 Configure the ISE or RADIUS server as the authorization server in the remote access VPN policy.

Step 9 Save and deploy the remote access VPN policy.

Related Topics

Configure Connection Profile Settings, on page 1832

Restrict Connection Profile Selection for a User Group

When you want to enforce a single connection profile on a user or user group, you can choose to disable the connection profile so that the group alias or URLs are not available for the users to select when they connect using the AnyConnect VPN module of Cisco Secure Client.

For example, if your organization wants to use specific configurations for different VPN user groups such as mobile users, corporate-issued laptop users, or personal laptop users, you can configure connection a profile specific to each of these user groups and apply the appropriate connection profile when the user connects to the VPN.

The AnyConnect VPN module of Cisco Secure Client, by default, shows a list of the connection profiles (by connection profile name, alias, or alias URL) configured in Cloud-Delivered Firewall Management Center and deployed on Firewall Threat Defense. If custom connection profiles are not configured, AnyConnect VPN module of Cisco Secure Client shows the *DefaultWEBVPNGroup* connection profile. Use the following procedure to enforce a single connection profile for a user group.

Before you begin

- On your Cloud-Delivered Firewall Management Center web interface, configure remote access VPN using the remote access VPN policy wizard with Authentication Method as 'Client Certificate Only' or 'Client Certificate + AAA'. Choose the username fields from the certificate.
- Configure ISE or RADIUS server for authorization and associate the group policy with the authorization server.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access policy and click **Edit**.

Step 3 Select **Access Interfaces** and disable **Allow users to select connection profile while logging in**.

Step 4 Click **Advanced Certificate Maps**.

Step 5 Select **Use the configured rules to match a certificate to a Connection Profile**.

Step 6 Select the **Certificate Map Name** or click the **Add** icon to add a certificate rule.

Step 7 Select the **Connection Profile**, and click **Ok**.

With this configuration, when a user connects from the AnyConnect VPN module of Cisco Secure Client, the user will have the mapped connection profile and will be authenticated to use the VPN.

Related Topics

Configure a group policy object, on page 3062
Configure Connection Profile Settings, on page 1832

Update the Secure Client Profile for Remote Access VPN Clients

Secure Client Profile is an XML file that contains an administrator-defined end user requirements and authentication policies to be deployed on a VPN client system as part of Secure Client. It makes the preconfigured network profiles available to end users.

You can use the GUI-based Secure Client Profile Editor, an independent configuration tool, to create them. The standalone profile editor can be used to create a new or modify existing Secure Client profile. You can download the profile editor from Cisco Software Download Center.

See the Secure Client Profile Editor chapter in the appropriate release of the Cisco Secure Client (including AnyConnect) Administrator Guide for details.

Before you begin

- Ensure that you have configured remote access VPN using the Remote Access Policy wizard and deployed the configuration on Firewall Threat Defense device. See Create a New Remote Access VPN Policy, on page 1822.
- On your Cloud-Delivered Firewall Management Center web interface, go to **Objects** > **Object Management** > **VPN** > **Secure Client File** and add the new Secure Client image.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access VPN policy and click **Edit**.

Step 3 Select the connection profile that includes the client profile to be edited, and click **Edit**.

Step 4 Click **Edit Group Policy** > **Secure Client** > **Profiles**.

Step 5 Select the client profile XML file from the list or click **Add** to add a new client profile.

Step 6 Save the group policy, connection profile, and then the remote access VPN policy.

Step 7 Deploy the changes.
Changes to the client profile will be updated on the VPN clients when they connect to the remote access VPN gateway.

Related Topics

Configure a group policy object, on page 3062

RADIUS Dynamic Authorization

Secure Firewall Threat Defense has the capability to use RADIUS servers for user authorization of VPN remote access and firewall cut-through-proxy sessions using dynamic access control lists (ACLs) or ACL names per user. To implement dynamic ACLs for dynamic authorization or RADIUS Change of Authorization (RADIUS CoA), you must configure the RADIUS server to support them. When the user tries to authenticate, the RADIUS server sends a downloadable ACL or ACL name to the Firewall Threat Defense. Access to a given service is either permitted or denied by the ACL. Secure Firewall Threat Defense deletes the ACL when the authentication session expires.

Related Topics

Add a RADIUS server group, on page 2943
Interfaces, on page 2980
Configuring RADIUS Dynamic Authorization, on page 1899
RADIUS Server Attributes for Secure Firewall Threat Defense, on page 1841

Configuring RADIUS Dynamic Authorization

Before you begin:

- Only one interface can be configured in the security zone or interface group if it is referred in a RADIUS Server.
- A dynamic authorization enabled RADIUS server requires Secure Firewall Threat Defense 6.3 or later for the dynamic authorization to work.
- Interface selection in RADIUS server is not supported on Secure Firewall Threat Defense 6.2.3 or earlier versions. The interface option will be ignored during deployment.
- Firewall Threat Defense posture VPN does not support group policy change through dynamic authorization or RADIUS change of authorization (CoA).

Table 111: Procedure

	Do This	More Info
Step 1	Log in to your Cloud-Delivered Firewall Management Center web interface.	
Step 2	Configure a RADIUS server object with dynamic authorization.	RADIUS server group configuration options, on page 2944
Step 3	Configure a route to ISE server through an interface enabled for change of authorization (CoA) to establish connectivity from Firewall Threat Defense to RADIUS server through routing or a specific interface.	RADIUS server group configuration options, on page 2944
Step 4	Configure a remote access VPN policy and select the RADIUS server group object that you have created with dynamic authorization.	Create a New Remote Access VPN Policy, on page 1822

	Do This	More Info
Step 5	Configure the DNS server details and domain-lookup interfaces using the Platform Settings.	Configure DNS, on page 1827 DNS server groups, on page 2965
Step 6	Configure split tunnelling in the remote access VPN group policy to allow DNS traffic through the VPN tunnel.	Configure a group policy object, on page 3062
Step 7	Deploy the configuration changes.	Deploy configuration changes, on page 202

Two-Factor Authentication

You can configure two-factor authentication for the remote access VPN. With two-factor authentication, the user must supply a username and static password, plus an additional item such as an RSA token or a passcode. Two-factor authentication differs from using a second authentication source in that two-factor is configured on a single authentication source, with the relationship to the RSA server tied to the primary authentication source.

Secure Firewall Threat Defense supports RSA tokens and Duo Push authentication requests to Duo Mobile for the second factor in conjunction with any RADIUS or AD server as the first factor in the two-factor authentication process.

Configuring RSA Two-Factor Authentication

About this task:

You can configure the RADIUS or AD server as the authentication agent in the RSA server, and use the server in Cloud-Delivered Firewall Management Center as the primary authentication source in the remote access VPN.

When using this approach, the user must authenticate using a username that is configured in the RADIUS or AD server, and concatenate the password with the one-time temporary RSA token, separating the password and token with a comma: *password*,*token*.

In this configuration, it is typical to use a separate RADIUS server (such as one supplied in Cisco ISE) to provide authorization services. You would configure the second RADIUS server as the authorization and, optionally, accounting server.

Before you begin:

Ensure that the following configurations are complete before configuring RADIUS two-factor authentication on Secure Firewall Threat Defense:

On the RSA Server

- Configure RADIUS or Active Directory server as an authentication agent.
- Generate and download the configuration (*sdconf.rec*) file.
- Create a token profile, assign the token to the user, and distribute the token to the user. Download and install the token on the remote access VPN client system.

For more information, see RSA SecureID Suite documentation.

On the ISE Server

- Import the configuration (*sdconf.rec*) file generated on the RSA server.
- Add the RSA server as the external identity source and specify the shared secret.

Table 112: Procedure

	Do This	More Info
Step 1	Log on to your Cloud-Delivered Firewall Management Center web interface.	
Step 2	Create a RADIUS server group.	RADIUS server group configuration options, on page 2944
Step 3	Create a RADIUS Server object within the new RADIUS server group, with RADIUS or AD server as the host and with a timeout of 60 seconds or more.	RADIUS server group configuration options, on page 2944 **Note** The RADIUS or AD server must be the same server that is configured as the authentication agent in RSA server. For two-factor authentication, make sure that the timeout is updated to 60 seconds or more in the Secure Client Profile XML file as well.
Step 4	Configure a new remote access VPN policy using the wizard or edit an existing remote access VPN policy.	Create a New Remote Access VPN Policy, on page 1822
Step 5	Select RADIUS as the authentication server and then select the newly-created RADIUS server group as the authentication server.	Configure AAA Settings for Remote Access VPN, on page 1835
Step 7	Deploy the configuration changes.	Deploy configuration changes, on page 202

Configuring Duo Two-Factor Authentication

About this task:

You can configure the Duo RADIUS server as the primary authentication source. This approach uses the Duo RADIUS Authentication Proxy. (You cannot use a direct connection with the Duo Cloud Service over LDAPS.)

For the detailed steps to configure Duo, see https://duo.com/docs/cisco-firepower.

You would then configure Duo to forward authentication requests directed to the proxy server to use another RADIUS server, or an AD server, as the first authentication factor, and the Duo Cloud Service as the second factor.

When using this approach, the user must authenticate using a username that is configured on both the Duo Cloud or web server, and the associated RADIUS server. The user must enter the password configured in the RADIUS server, followed by one of the following Duo codes:

- **Duo-passcode**. For example, *my-password*,123456.
- **push**. For example, *my-password*,**push**. Use push to tell Duo to send a push authentication to the Duo Mobile app, which the user must have already installed and registered.
- **sms**. For example, *my-password*,**sms**. Use **sms** to tell Duo to send an SMS message with a new batch of passcodes to the user's mobile device. The user's authentication attempt will fail when using **sms**. The user must then re-authenticate and enter the new passcode as the secondary factor.
- **phone**. For example, *my-password*,**phone**. Use **phone** to authenticate using phone callback.

For more information on login options with examples, see https://guide.duo.com/anyconnect.

Before you begin:

Before configuring two-factor authentication with Duo Authentication Proxy on Firewall Threat Defense, ensure that you complete the following configurations:

- Configure a working primary authentication (RADIUS or AD) for your remote access VPN users before you begin to deploy Duo.
- Install Duo proxy service on a Windows or Linux machine within your network to integrate Duo with Secure Firewall Threat Defense remote access VPN. This Duo proxy server also acts as a RADIUS server.

  Download and install the most recent Duo authentication proxy from the following location:

  - **Windows**: https://dl.duosecurity.com/duoauthproxy-latest.exe
  - **Linux**: https://dl.duosecurity.com/duoauthproxy-latest-src.tgz
  - Verify the checksum at https://duo.com/docs/checksums#duo-authentication-proxy.

- Configure Duo authentication file `authproxy.cfg`. Follow instructions on the https://duo.com/docs/cisco-firepower#configure-the-proxy page to configure the authentication configuration settings.

  The `authproxy.cfg` configuration file must contain the details for RADIUS or ISE server, Firewall Threat Defense device, Duo proxy server details, Integration Key, Secret key, and API host details.
- Ensure that you have the right API host information in the `authproxy.cfg` file.
- Configure other required settings such as secondary authentication factor in the newly installed Duo proxy server at **Duo Security Server** > **Duo Admin Panel** > **Applications** > **CISCO RADIUS VPN**.

Table 113: Procedure

	Do This	More Info
Step 1	Log on to your Cloud-Delivered Firewall Management Center web interface.	
Step 2	Create a RADIUS server group.	RADIUS server group configuration options, on page 2944

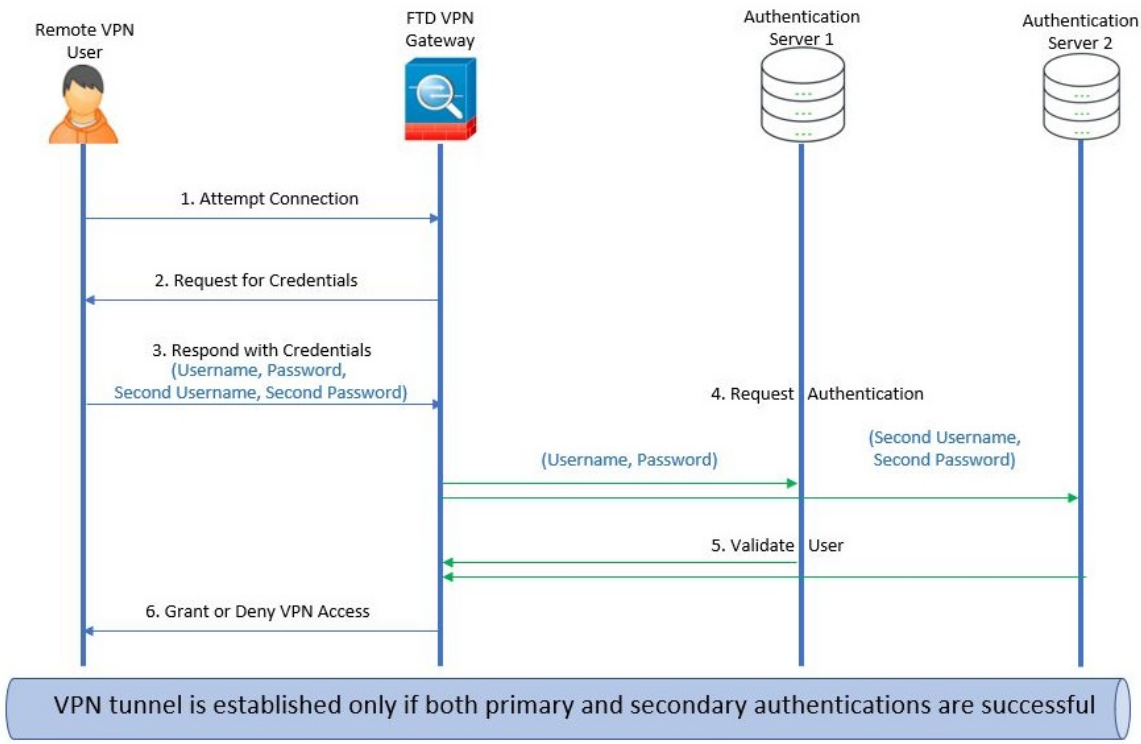
	Do This	More Info
Step 3	Create a RADIUS Server object within the new RADIUS server group with Duo proxy server as the host with a timeout of 60 seconds or more.	Configure RADIUS server options, on page 2945 **Note** For two-factor authentication, make sure that the timeout is updated to 60 seconds or more in the Secure Client Profile XML file as well.
Step 4	Configure a new remote access VPN policy using the wizard or edit an existing remote access VPN policy.	Create a New Remote Access VPN Policy, on page 1822
Step 5	Select RADIUS as the authentication server and then select the RADIUS server group created with the Duo proxy server as the authentication server.	Configure AAA Settings for Remote Access VPN, on page 1835
Step 7	Deploy the configuration changes.	Deploy configuration changes, on page 202

Secondary Authentication

Secondary authentication or double authentication in Secure Firewall Threat Defense adds an additional layer of security to remote access VPN connections by using two different authentication servers. With secondary authentication enabled, the Secure Client VPN users must provide two sets of credentials to login to the VPN gateway.

Secure Firewall Threat Defense remote access VPN supports secondary authentication in AAA Only and Client Certificate & AAA authentication methods.

Figure 375: Remote Access VPN Secondary or Double Authentication



Related Topics

Configure Remote Access VPN Secondary Authentication, on page 1904

Configure Remote Access VPN Secondary Authentication

When remote access VPN authentication is configured to use both client certificate and authentication sever, VPN client authentication is done using both the client certificate validation and AAA server.

Before you begin

- Configure two authentication (AAA) servers— the primary and secondary authentication servers, and required identity certificates. The authentication servers can be RADIUS server, and AD or LDAP realms.
- Ensure that the AAA servers are reachable from the Secure Firewall Threat Defense device for the remote access VPN configuration to work. Configure routing (at **Devices** > **Device Management**, click **Edit** (). Then, click **Device > Routing**) to ensure connectivity to the AAA servers.

Procedure

Step 1 On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access** .

Step 2 Select a remote access policy and click **Edit**; or click **Add** to create a new remote access VPN policy.

Step 3 For a new remote access VPN policy, configure the authentication while selecting connection profile settings. For an existing configuration, select the connection profile that includes the client profile, and click **Edit**.

Step 4 Click **AAA** > **Authentication Method**, **AAA** or **Client Certificate & AAA**.

- When you select the **Authentication Method** as:

  Client Certificate & AAA—Authentication is done using both client certificate and AAA server.

  - **AAA**—If you select the **Authentication Server** as **RADIUS**, by default, the Authorization Server has the same value. Select the **Accounting Server** from the drop-down list. Whenever you select **AD** and **LDAP** from the Authentication Server drop-down list, you must manually select the **Authorization Server** and **Accounting Server** respectively.
  - Whichever authentication method you choose, select or deselect **Allow connection only if user exists in authorization database**.

- **Use secondary authentication** — Secondary authentication is configured in addition to primary authentication to provide additional security for VPN sessions. Secondary authentication is applicable only to **AAA only** and **Client Certificate & AAA** authentication methods.

  Secondary authentication is an optional feature that requires a VPN user to enter two sets of username and password on the Secure Client login screen. You can also configure to pre-fill the secondary username from the authentication server or client certificate. Remote access VPN authentication is granted only if both primary and secondary authentications are successful. VPN authentication is denied if any one of the authentication servers is not reachable or one authentication fails.

  You must configure a secondary authentication server group (AAA server) for the second username and password before configuring secondary authentication. For example, you can set the primary authentication server to an LDAP or Active Directory realm and the secondary authentication to a RADIUS server.

  Note
  By default, secondary authentication is not required.

  Authentication Server— Secondary authentication server to provide secondary username and password for VPN users.

  Select the following under **Username for secondary authentication**:

  - **Prompt**: Prompts the users to enter the username and password while logging on to VPN gateway.
  - **Use primary authentication username**: The username is taken from the primary authentication server for both primary and secondary authentication; you must enter two passwords.
  - **Map username from client certificate**: Prefills the secondary username from the client certificate.
    - If you select **Map specific field** option, which includes the username from the client certificate. The **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively. If you select the **Use entire DN (Distinguished Name) as username** option, the system automatically retrieves the user identity.

      See **Authentication Method** descriptions for more information about primary and secondary field mapping.
    - **Prefill username from certificate on user login window**: Prefills the secondary username from the client certificate when the user connects via Secure Client.
      - **Hide username in login window**: The secondary username is pre-filled from the client certificate, but hidden to the user so that the user does not modify the pre-filled username.

- **Use secondary username for VPN session**: The secondary username is used for reporting user activity during a VPN session.

For more information, see Configure AAA Settings for Remote Access VPN, on page 1835.

Related Topics

Configure Connection Profile Settings, on page 1832

Single Sign-On Authentication with SAML 2.0

About SAML Single Sign-On Authentication

Security Assertion Markup Language (SAML) is an open standard for logging users into applications using their sessions in another context. Organizations already know the identity of users when users log in to their Active Directory (AD) domain or the intranet. They use this identity information to log in users to other applications, such as web-based applications using SAML. Individual applications do not need to store credentials and users do not have to remember and manage different sets of credentials for individual applications. SAML single sign-on (SSO) works by transferring the user's identity from one place (the identity provider) to another (the service provider).

SAML Single Sign-On with Secure Firewall Threat Defense

The Secure Firewall Threat Defense device supports SAML 2.0 single sign-on (SSO) authentication for remote access VPN connections using the Secure Client. You need the following to configure SAML 2.0 SSO on Secure Firewall Threat Defense:

- **Identity Provider (IdP)**—The Duo Access Gateway acts as the identity provider to perform user authentication and issues assertions.
- **Service Provider (SP)**—The Firewall Threat Defense device acts as the service provider and obtains the authentication assertion from the identity provider.
- **VPN Client**—The Secure Client performs SAML 2.0 authentication via the embedded browser.

Guidelines and Limitations for SAML 2.0

- Firewall Threat Defense supports the following signatures for SAML authentication:
  - SHA1 with RSA and HMAC
  - SHA2 with RSA and HMAC
- Firewall Threat Defense supports SAML 2.0 Redirect-POST binding , which is supported by all SAML IdPs.
- The Firewall Threat Defense functions as a SAML SP only. It cannot act as an Identity Provider in gateway mode or peer mode.
- You can enforce an access policy on a SAML-authenticated user if you have an associated identity policy with an AD realm matching the SAML domain. However, it does not work for Azure AD SAML because it requires additional mapping from the tenant ID of the Azure AD to an associated realm ID on the threat defense device.

- Having SAML authentication attributes available in DAP evaluation (similar to RADIUS attributes sent in RADIUS auth response from AAA server) is not supported. Firewall Threat Defense supports SAML enabled group policy on DAP policy; however, you cannot check the username attribute while using SAML authentication, because the username attribute is masked by the SAML Identity provider.
- Firewall Threat Defense administrators need to ensure clock synchronization between the Firewall Threat Defense and the SAML IdP for proper handling of authentication assertions and proper timeout behavior.
- Firewall Threat Defense administrators have the responsibility to maintain a valid signing certificate on both Firewall Threat Defense and IdP considering the following:
  - The IdP signing certificate is mandatory when configuring an IdP on the Firewall Threat Defense.
  - The Firewall Threat Defense does not do a revocation check on the signing certificate received from the IdP.
- In SAML assertions, there are NotBefore and NotOnOrAfter conditions. The Firewall Threat Defense SAML configured timeout interacts with these conditions as follows:
  - Timeout overrides NotOnOrAfter if the sum of NotBefore and timeout is earlier than NotOnOrAfter.
  - If NotBefore + timeout is later than NotOnOrAfter, then NotOnOrAfter takes effect.
  - If the NotBefore attribute is absent, the Firewall Threat Defense denies the login request. If the NotOnOrAfter attribute is absent and SAML timeout is not set, Firewall Threat Defense denies the login request.
- Firewall Threat Defense does not work with Duo in a deployment using an internal SAML, which forces the Firewall Threat Defense to proxy for the client to authenticate, due to the FQDN change that occurs during challenge/response for Two-factor authentication (push, code, password).
- When using SAML with Secure Client, follow these guidelines:
  - Untrusted server certificates are not allowed in the embedded browser.
  - The embedded browser SAML integration is not supported in CLI or SBL modes.
  - SAML authentication established in a web browser is not shared with Secure Client and vice versa.
  - Depending on the configuration, various methods are used when connecting to the headend with the embedded browser. For example, while Secure Client might prefer an IPv4 connection over an IPv6 connection, the embedded browser might prefer IPv6, or vice versa. Similarly, Secure Client may fall back to no proxy after trying proxy and getting a failure, while the embedded browser may stop navigation after trying proxy and getting a failure.
  - You must synchronize your Firewall Threat Defense's Network Time Protocol (NTP) server with the IdP NTP server in order to use the SAML feature.
  - You cannot access internal servers with SSO after logging in using an internal IdP.
  - The SAML IdP NameID attribute determines the user's username and is used for authorization, accounting, and VPN session database.
- SAML does not support Start Before Logon (SBL).
- Multiple attributes received with a SAML assertion are not supported.

- A firewall uses the Identity Provider Entity ID as the SAML object name from the single sign-on server object created in the Cloud-Delivered Firewall Management Center (**Objects** > **Object Management** > **AAA Server** > **Single Sign-on Server**). Hence, you cannot use multiple SAML objects with the same Identity Provider Entity ID on a single firewall.

Configuring a SAML Single Sign-On Authentication

Before you begin

Ensure that you have done the following before you configure SAML single sign-on with Firewall Threat Defense remote access VPN:

- Create an account with Duo.
- Download and install the Duo Access Gateway.
- Obtain the following from your SAML identity provider (Duo).
  - Identity Provider Entity ID URL
  - Sign-in URL
  - Sign-out URL
  - Identity provider certificate
- Create a SAML single sign-on server object. For more information, see Add a single sign-on server, on page 2946.

Note You can create a single sign-on server object in the **Connection Profile** settings when you create a new policy using the Remote Access VPN policy Wizard.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** next to the remote access VPN policy for which you want to configure SAML authentication. If you want to create a new policy, click **Add**.

Step 3 Click **Edit** on the connection profile that you want to modify.

Step 4 Choose **AAA** settings and select **SAML** from the **Authentication Method** drop-down.

Step 5 Choose the required SAML single sign-on server as the **Authentication Server**.

Step 6 Configure the required settings for the remote access VPN.

Step 7 Save and deploy the remote access VPN policy on your Firewall Threat Defense device.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1835

Configuring SAML Authorization

About SAML Authorization

SAML authorization supports user attributes delivered in SAML assertions within the AAA and Dynamic Access Policy (DAP) frameworks. You can configure the SAML assertion attributes on the Identity Provider as name-value pairs which then parses as strings. The attributes received are made available to DAP so that they can be used when defining selection criteria within a DAP record. The SAML assertion *cisco_group_policy* is used to determine the Group Policy to be applied to the VPN session.

Dynamic Access Policy Attribute Representation

In the DAP table, the DAP attributes are represented in the following format:

```
aaa.saml.name = "value”
```

Example, *aaa.saml.department = ”finance"*

This attribute can be used in DAP selection as follows:

```
<attr>
<name>aaa.saml.department</name>
<value>finance</value>
<operation>EQ</operation>
</attr>
```

Multi-Valued Attributes

Multi-valued attributes are also supported in DAP and the DAP table is indexed :

```
aaa.saml.name.1 = "value”
aaa.saml.name.2 = "value"
```

Active Directory memberOf Attributes

The Active Directory (AD) memberOf attribute receives a special processing that is consistent with the way it is handled through an LDAP query.

Group names are represented by the CN attribute of the DN.

Example Attributes received from the authorization server:

```
memberOf = "CN=FTD-VPN-Group,OU=Users,OU=TechspotUsers,DC=techspot,DC=us"
 memberOf = "CN=Domain Admins,OU=Users,DC=techspot,DC=us”
```

Dynamic Access Policy attributes:

```
aaa.saml.memberOf.1 = "FTD-VPN-Group”
aaa.saml.memberOf.2 = "Domain Admins"
```

Interpretation of the cisco_group_policy Attribute

A group-policy can be specified by a SAML assertion attribute. When an attribute "cisco_group_policy" is received by the Firewall Threat Defense, the corresponding value is used to select the connection group-policy

Configure SAML Authorization

Before you begin

Ensure that you have configured a single sign-on server like DUO and completed the required Identity Provider (IdP) and Service Provider (SP) settings.

For more information, see Single Sign-On Authentication with SAML 2.0, on page 1906.

Procedure

Step 1 Configure a single sign-on server object if not configured already.

a) Choose **Objects** > **Object Management** > **AAA Server** > **Single Sign-on Server**.
b) Click **Add Single Sign-on Server**.
c) Enter the single sign-on server details and click **Save**.

For more information, see Add a single sign-on server, on page 2946.

Step 2 Configure SAML authentication in the remote access VPN connection profile.

a) Choose **Devices** > **VPN** > **Remote Access**.
b) Click **Edit** on the remote access VPN policy for which you want to configure SAML authorization or create a new policy.
c) Edit the required connection profile and select **AAA**.
d) Select the single sign-on server object from the **Authentication Server** drop-down.
e) Save the remote access VPN configuration.

Step 3 Match a SAML criteria in DAP policy.

a) Select **Devices** > **VPN** > **Dynamic Access Policy**.
b) Create a new DAP or edit an existing one.
c) Create a DAP record or edit and existing record.
d) Click **AAA Criteria** > **SAML Criteria** > **Add SAML Criteria**.
e) Create a SAML criteria based on the SAML assertions returned by the SSO server.

Step 4 Deploy the remote access VPN configuration.

Related Topics

Configure Connection Profile Settings, on page 1832
Firewall Threat Defense group policy objects, on page 3061

Advanced Secure Client Configurations

Configure Secure Client Modules on a Firewall Threat Defense

Secure Client can integrate with various Cisco endpoint security solutions and offer enhanced security using different Secure Client modules.

You can use the managed headend Firewall Threat Defense to distribute and manage Secure Client modules to the endpoints. When a user connects to the Firewall Threat Defense, it downloads and installs Secure Client and the required modules on the endpoint.

In version 6.7 and later, you can use the headend Firewall Threat Defense, managed by a Cloud-Delivered Firewall Management Center, to distribute and manage Secure Client modules to the endpoints. These modules then integrate with the corresponding Cisco endpoint security solution.

In versions 6.4 to 6.6, you can enable these modules and profiles on a Firewall Threat Defense using FlexConfig. For more information, see Configure AnyConnect Modules and Profiles Using FlexConfig.

Benefits

If you use a Firewall Threat Defense to distribute and manage Secure Client modules to the endpoints, you can easily perform the following tasks:

- Distribute and manage Secure Client modules and profiles on each endpoint.
- Upgrade Secure Client on each endpoint.

Types of Secure Client Modules

AMP Enabler

Use this module to deploy Secure Endpoint, on endpoints. The module pushes Secure Endpoint to endpoints from a server hosted locally within the enterprise. This module provides an additional security agent that detects potential malware threats in the network, removes these threats, and protects the enterprise.

In Cisco Secure Client 5.0, AMP Enabler is only for macOS. Cisco Secure Client for Windows offers full integration with Cisco Secure Endpoint.

ISE Posture

Use this module to perform endpoint posture checks such as antivirus, antispyware, operating system and so on using Cisco Identity Services Engine (ISE) and assess the endpoint's compliance. ISE provides next generation identity and access control policy. ISE Posture performs a client-side evaluation. The client receives the posture requirement policy from the headend, performs the posture data collection, compares the results against the policy, and sends the assessment results back to the headend.

Network Visibility

Use this module to monitor the endpoint application usage using the Network visibility module. You can uncover potential behavior anomalies and make informed network design decisions. It enhances the enterprise administrator's ability to do capacity and service planning, auditing, compliance, and security analytics. You can share the usage data with NetFlow analysis tools such as Cisco Stealthwatch.

Umbrella Roaming Security

Use this module for a DNS-layer security using the Cisco Umbrella Roaming Security service. Cisco Umbrella provides content filtering, multiple policies, robust reporting, active directory integration, and much more.

Web Security

Use this module to enable Cisco Secure Web Appliance (SWA), powered by Cisco Talos. This module protects the endpoint by blocking risky sites and testing unknown sites before allowing users to access them. It can deploy web security either through the on-prem WSA or the cloud-based Cisco Cloud Web Security. This module is not part of the AnyConnect package from release 4.5 and in Secure Client 5.0.

Network Access Manager

This module provides a secure layer 2 network and performs device authentication to access wired and wireless networks. Network Access Manager manages user and device identity and the network access protocols required for secure access.

Network Access Manager is not supported on macOS or Linux.

Start Before Login

Start Before Login (SBL) allows users to establish their VPN connection to the enterprise infrastructure before logging onto Windows. After the SBL module installation, you must enable SBL in the Secure Client VPN profile and add it to the remote access VPN group policy.

DART

Diagnostics and Reporting Tool (DART) collates system logs and other diagnostic information to troubleshoot AnyConnect installation and connection problems. You can send this data to Cisco TAC for troubleshooting.

By default, DART is not enabled in new RA VPN group policies for 6.7 and later versions. In 6.6 and earlier versions, DART is enabled by default.

Feedback

The customer experience feedback (CEF) module provides information about which features and modules you use and have enabled. This information gives an insight into the user experience so that Cisco can continue to improve the quality, reliability, performance, and user experience of the Cisco Secure Client. Secure Client does not download the Feedback module to the endpoint. The feedback data is sent to the Cisco Feedback Server.

Prerequisites for Configuring Secure Client Modules

- Configure the associated products depending on the module that you are going to use.
- Download the following Secure Client-related packages from Cisco Software Download Center to your local host.
  - Cisco Secure Client Headend Deployment Package for the required platforms.

    This package is for the headend and contains all the Secure Client modules. For Windows, the filename is cisco-secure-client-win-5.0.03076-webdeploy-k9.pkg.
  - Profile Editor: Create profiles for the modules that require profiles.

    Secure Client needs a Secure Client profile for some of the modules. A profile contains configurations to enable the modules and connect to the corresponding security services. The profile editor supports only Windows.

    The following table lists if the modules require a client profile:

Secure Client Module	Requires a Client Profile
AMP Enabler	Yes
ISE Posture	Yes
Network Access Manager	Yes
Network Visibility Module	Yes
Umbrella Roaming Secure Module	Yes
Feedback	Yes
Web Security	Yes
DART	No
Start Before Login	No

- Licensing
  - You need one of the following Secure Client licenses: Secure Client Premier, Secure Client Advantage, or Secure Client VPN Only
  - Your Cloud-Delivered Firewall Management Center Essentials license must allow export-controlled functionality.

    Choose **System** (⚙) > **Licenses** > **Smart Licenses** to verify this functionality in the management center.

Guidelines for Configuring Secure Client Modules

- All Secure Client modules are supported from AnyConnect 4.8 and later, and Secure Client 5.0.
- Different modules support profiles with different file extensions. The following table lists the modules and the supported file extensions of their profiles:

Table 114: Supported File Extensions of Profiles

Module	File Extension
AMP Enabler	*.xml, *.asp
Feedback	*.xml
ISE Posture	*.xml, *.isp
Network Access Manager	*.xml, *.nsp
Network Visibility	*.xml, *.nvmsp
Umbrella Roaming Security	*.xml, *.json
Web Security	*.xml, *.wsp, *.wso

- You can add only one entry per client module. You can edit or delete an entry for a module.
- If you plan to use ISE posture and Network Access Manager modules on a Windows OS, you must install Network Access Manager before you use the ISE Posture module.
- If you enable the Umbrella Roaming Security module, ensure that you disable the **Always send DNS requests over tunnel** option under split tunnelling in the VPN group policy.
- If you want to use SBL, then you must enable SBL in the Secure Client VPN profile.

Install Secure Client Modules using a Firewall Threat Defense

Before you begin

Ensure that you review the Prerequisites for Configuring Secure Client Modules, on page 1912 and Guidelines for Configuring Secure Client Modules, on page 1913 topics.

Procedure

Step 1 The administrator creates profiles, if needed, for the required Secure Client modules.

Step 2 The administrator uses the Cloud-Delivered Firewall Management Center to:

a) Configure the modules and add the profiles in the remote access VPN group policy.

b) Deploy the configuration on the Firewall Threat Defense.

Step 3 The user uses Secure Client to initiate a VPN connection to the Firewall Threat Defense.

Step 4 The Firewall Threat Defense authenticates the user.

Step 5 The Secure Client checks for updates.

Step 6 The Firewall Threat Defense distributes the Secure Client modules and the profiles on the endpoint.

What to do next

Configure a Remote Access VPN Group Policy with Secure Client Modules, on page 1914.

Configure a Remote Access VPN Group Policy with Secure Client Modules

To install and update the Secure Client modules on the endpoint using a Firewall Threat Defense managed by a Cloud-Delivered Firewall Management Center, you must update the remote access VPN group policy with the Secure Client module configurations.

Before you begin

Ensure that you have configured a remote access VPN policy in the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.
Step 2 Select a remote access VPN policy and click **Edit**.
Step 3 Select a connection profile and click **Edit**.
Step 4 Click **Edit Group Policy**.
Step 5 Click the **Secure Client** tab.
Step 6 Click **Client Modules**.
Step 7 Click +.
Step 8 Choose a module from the **Client Module** drop-down list.
Step 9 Choose a profile for the module from the **Profile to download** drop-down list or click + to add a profile.
Step 10 Check the **Enable module download** check box to download the module on the endpoint.
Step 11 Click **Add**.
Step 12 Repeat steps 7 to 11 if you want to add more modules.
Step 13 Click **Save**.

What to do next

1. Deploy the configuration on the threat defense.
2. Launch the Secure Client, select the VPN profile, and connect to the VPN. Secure Client installs the configured modules on it.
3. Verify the configuration. For more information, see Verify Secure Client Modules Configuration, on page 1915.

Verify Secure Client Modules Configuration

On the Firewall Threat Defense

Use the following commands on the Firewall Threat Defense to view the profiles and the Secure Client modules configuration:

- **show disk0:**- View the profiles and their configuration.
- **show run webvpn** - View details of the Secure Client configurations.
- **show run group-policy <ravpn_group_policy_name>** - View details of the RA VPN group policy for Secure Client.
- **show vpn-sessiondb anyconnect** - View details of the active Secure Client VPN sessions.

On the Endpoint

1. Use the Secure Client to establish a VPN connection to the Firewall Threat Defense.

2. Verify if the configured modules are downloaded and installed as part of the Secure Client.
3. Verify if the configured profiles, if any, are available in the locations documented in Profile Locations for all Operating Systems.

On the Cloud-Delivered Firewall Management Center

You can monitor remote access VPN active sessions on the Cloud-Delivered Firewall Management Center using the Remote Access VPN dashboard (**Overview** > **Dashboards** > **Remote Access VPN**). You can quickly determine problems related to user sessions and mitigate the problems for your network and users.

Configure Application-Based (Per App VPN) Remote Access VPN on Mobile Devices

When you use Secure Client to establish a VPN connection from a mobile device, all the traffic including the traffic from personal applications is routed through the VPN.

For mobile devices that run on Android or iOS, you can restrict the applications that use the VPN tunnel. This application-based remote access VPN is called Per App VPN. To use Per App VPN, you must install and configure a third-party Mobile Device Manager (MDM) application. You must define the list of approved applications that can be used over the VPN tunnel in the MDM. You can enable Per App VPN on the Firewall Threat Defense headend so that your MDM can apply your policies on mobile devices.

Benefits

Benefits of restricting the remote access VPN to approved applications include:

- Performance—Limits VPN traffic over the corporate network and frees up resources of the VPN headend.
- Protection—Protects the corporate VPN tunnel from unapproved malicious applications on the mobile device.

Prerequisites and Licensing for Configuring Per App VPN Tunnels

Prerequisites

- Install and configure a third-party Mobile Device Manager (MDM).

  You must configure the applications that will be allowed in the VPN in the MDM itself, not on the Firewall Threat Defense headend device.
- Download the Cisco AnyConnect Enterprise Application Selector from Cisco Software Download Center.

  You need this tool to define the Per App VPN policy.

Licensing

- Secure Client Premier, or Secure Client Advantage.
- Essentials license must allow export-controlled functionality.

  To verify this functionality in the Cloud-Delivered Firewall Management Center, choose **System** (⚙) > **Licenses** > **Specific Licenses**.

Determine the Application IDs for Mobile Applications

Before configuring the Firewall Threat Defense headend to allow application-based VPN from mobile devices, you must determine which applications should be allowed in the tunnel.

We strongly recommend that you configure the per-app policy in the MDM on the user's mobile device. This simplifies the headend configuration. If you decide to configure the list of allowed applications on the headend, you must determine the application IDs for each application on each type of endpoint.

The application ID, called the bundle ID in iOS, is a reverse DNS name. You can use an asterisk as a wildcard. For example, ***.*** indicates all applications, com.cisco.* indicates all Cisco applications.

To determine the application IDs:

- **Android**—Go to Google Play in a web browser and select the Apps category. Click (or hover over) an application that you want to allow, then look at the URL. The app id is in the URL, on the **id=** parameter. For example, the following URL is for Facebook Messenger, so the app id is com.facebook.orca.

  https://play.google.com/store/apps/details?id=com.facebook.orca

  For applications that are not available through Google Play, such as your own applications, download a package name viewer application to extract the app ID. There are many of these applications available, one of them should provide what you need, but Cisco does not endorse any of them.

- **iOS**—There is no straight-forward way to get the bundle ID. Following is one way to find it:

  1. Use a desktop web browser such as Chrome to search for the application name.
  2. In the search results, look for the link to download the app from the Apple App Store. For example, Facebook Messager would be similar to:

     https://apps.apple.com/us/app/messenger/id454638411
  3. Copy the number after the **id** string. In this example, **454638411**.
  4. Open a new browser window, and add the number to the end of the following URL:

     https://itunes.apple.com/lookup?id=

     For this example: https://itunes.apple.com/lookup?id=454638411
  5. You will be prompted to download a text file, usually named 1.txt. Download the file.
  6. Open the file in a text editor such as WordPad, and search for bundleId. For example:

     "bundleId":"com.facebook.Messenger",

     In this example, the bundle ID is com.facebook.Messenger. Use this as the app ID.

Once you have your list of application IDs, you can configure the policy as explained in .

Configure Application-Based VPN Tunnels

After you install and configure your MDM software, you can enable Per App VPN on the Firewall Threat Defense headend device. Once enabled on the headend, your MDM software will control which applications are tunneled over the VPN to the corporate network.

Before you begin

- Ensure that you have a remote access VPN policy in the Cloud-Delivered Firewall Management Center.

- Configure Per App VPN using MDM and enroll each device to the MDM server.
- Download the Cisco AnyConnect Enterprise Application Selector.

Procedure

Step 1 Use the Cisco AnyConnect Enterprise Application Selector to define the Per App VPN policy.

We recommend that you create a simple **Allow All** policy, and define the allowed applications in the MDM. However, you can specify a list of applications to allow and control the list from the headend. If you want to include specific applications, create a separate rule for each application, using a unique name and the application's app ID. For more information on getting the app IDs, see Determine the Application IDs for Mobile Apps.

To create an **Allow All** policy that supports both Android and iOS platforms using the AnyConnect Enterprise Application Selector:

a) Choose **Android** from the drop-down list as the platform type and configure the following options:

- **Friendly Name**—Enter a name for the policy. For example, Allow_All.
- **App ID**—Enter *.* to match all possible applications.
- Leave the other options.

b) Choose **iOS** from the drop-down list as the platform type and configure the following options:

- **Friendly Name**—Enter a name for the policy. For example, Allow_All.
- **App ID**—Enter *.* to match all possible applications.
- Leave the other options.

c) Choose **Policy > View Policy** to get the base64 encoded string for the policy.

This string contains an encrypted XML file that allows the Firewall Threat Defense to see the policies. Copy this value. You need this string when you configure Per App VPN on the Firewall Threat Defense.

Step 2 Use the Cloud-Delivered Firewall Management Center to enable the Per App on the Firewall Threat Defense headend device.

a) Choose **Devices** > **VPN** > **Remote Access**.
b) Select a remote access VPN policy and click **Edit**.
c) Select a connection profile and click **Edit**.
d) Click **Edit Group Policy**.
e) Click the **Secure Client** tab.
f) Click **Custom Attributes** and click **+**.
g) Choose **Per App VPN** from the **Secure Client Attribute** drop-down list.
h) Choose an object from the **Custom Attribute Object** drop-down list or click **+** to add an object.

When you add a new custom attribute object for Per App VPN, enter the name, description, and the base64 encoded policy string from the Cisco AnyConnect Enterprise Application Selector.

i) Click **Save**.

j) Click **Add** and click **Save**.

Step 3 Deploy your changes on the Cloud-Delivered Firewall Management Center.

What to do next

1. Launch the Secure Client, select the VPN profile, and connect to the VPN.
2. Verify the configuration. For more information, see Verify Per App Configuration, on page 1919.

Verify Per App Configuration

On the Firewall Threat Defense

Use the following commands on the Firewall Threat Defense to view the Per App configuration:

- **show run webvpn**
- **show run group-policy <ravpn_group_policy_name>**
- **show run anyconnect-custom-data**

On the Endpoint

After the endpoint establishes a VPN connection with the Firewall Threat Defense:

1. Click the **Statistics** icon of the Secure Client.
2. **Tunnel Mode** will be “Application Tunnel” instead of “Tunnel All Traffic.”
3. **Tunneled Apps** will list the applications you enabled for tunneling in the MDM.

Remote Access VPN Examples

How to Limit Secure Client Bandwidth Per User

This section provides instructions to limit the maximum bandwidth that the VPN users consume when they connect using the Secure Client to Secure Firewall Threat Defense remote access VPN gateway. You can limit the maximum bandwidth by using a Quality of service (QoS) policy in Firewall Threat Defense, to ensure that a single user or group or users do not take over the entire resource. This configuration lets you give priority to critical traffic, prevent bandwidth hogging, and manage the network. If a When traffic exceeds the maximum rate, the Firewall Threat Defense drops the excess traffic.

Step	Do This	More Info
1	Create and set up a realm.	Create an LDAP realm or an Active Directory realm and realm directory, on page 2325

Step	Do This	More Info
2	Create a QoS policy and QoS rule for the user or group available in the newly created realm.	• See Create a QoS policy, on page 1696 to create a QoS policy. • See Configuring QoS rules, on page 1698 to create a QoS rule.
3	Configure remote access VPN policy and select the newly created realm for user authentication.	Create a New Remote Access VPN Policy, on page 1822
4	Deploy the remote access VPN policy.	Deploy configuration changes, on page 202

How to Use VPN Identity for User-Id Based Access Control Rules

Step	Do This	More Info
1	Create and set up a realm.	Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
2	Create an identity policy and add an identity rule.	• See Create an identity policy, on page 2507 to create an identity policy. • See Create an identity rule, on page 2515 to create an identity rule.
3	Associate the identity policy with an access control policy.	Associate prefilter, decryption, and identity policies with an access control, on page 2016
4	Configure remote access VPN policy and select the newly created realm for user authentication.	Create a New Remote Access VPN Policy, on page 1822
5	Deploy the remote access VPN policy.	Deploy configuration changes, on page 202

Configure Firewall Threat Defense Multiple Certificate Authentication

Multiple Certificate-based Authentication

Multiple certificate-based authentication allows the Firewall Threat Defense to validate the machine or device certificate. Multiple certificates can be enabled for certificate-based authentication in the remote access VPN connection profile. It can be combined with AAA authentication. The multiple certificates option in the remote access VPN connection profile allows certificate authentication of both the machine and user via certificates. This ensures that the device is a corporate-issued device, in addition to authenticating the user's identity certificate to allow RA VPN access. The administrator can choose if the username for the session should be taken from the machine certificate or user certificate.

When multiple certificate-based authentication is configured, two certificates are obtained from the VPN client:

- **First Certificate** —Machine certificate to authenticate the endpoint.
- **Second Certificate**—User certificate to authenticate the VPN user.

For detailed information about Firewall Threat Defense certificates, see Manage Firewall Threat Defense certificates, on page 3080.

Limitations

- Multiple certificate authentication currently limits the number of certificates to two.
- Secure Client supports only RSA-based certificates.
- Only SHA256, SHA384, and SHA512 based certificates are supported during the Secure Client aggregate authentication.
- Certificate authentication cannot be combined with SAML authentication.

Pre-fill Username from Certificate

The Pre-fill username option allows a field from the certificates to be parsed and used for subsequent AAA authentication (primary and secondary). When two certificates are used for authentication, the Administrator can choose the certificate from which the username should be derived for the prefill functionality. By default, username for prefill is retrieved from the User certificate (second certificate received from Secure Client). The prefilled username is used as the VPN session username when the Certificate Only authentication method is enabled. When AAA and certificate authentication is enabled, VPN session username will be based on the pre-fill option.

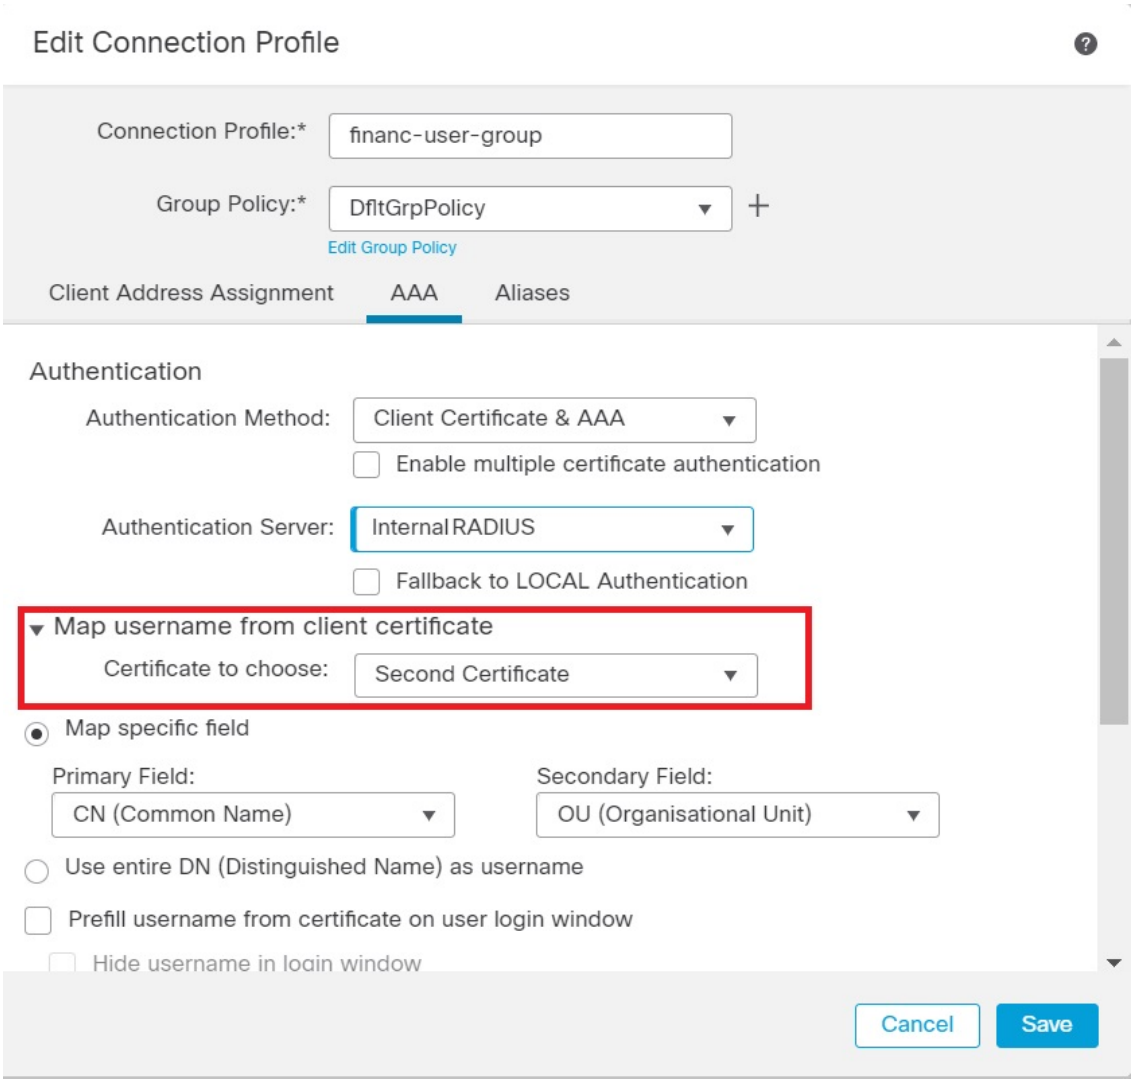
Configure Multiple Certificate Authentication for Remote Access VPN

1. On your Cloud-Delivered Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.
2. Edit an existing remote access policy, or create a new one and then edit.

   See Create a New Remote Access VPN Policy, on page 1822.
3. Select the connection profile to configure multiple certificate authentication, and click **Edit**.

   See Configure Connection Profile Settings, on page 1832.
4. Choose **AAA**, and then select an **Authentication Method**:

Figure 376:



- **Client Certificate Only**—User is authenticated using client certificate. Client certificate must be configured on VPN client endpoints. By default, the username is derived from client certificate fields CN & OU respectively. In case, the username is specified in other fields in the client certificate, use 'Primary' and 'Secondary' field to map appropriate fields.
- **Client Certificate & AAA**—User is authenticated using both the types of authentication, AAA and client certificate.

5. Select **Enable multiple certificate authentication**.
6. Select **Map username from client certificate** and select a certificate from the **Certificate to choose** drop-down to choose the username for the VPN session from the machine certificate or user certificate.
   - **First Certificate** —Map the username from the Machine Certificate.
   - **Second Certificate**—Map the username from the User certificate to authenticate the VPN user.

7. Configure the required connection profile settings and remote access VPN settings.
8. Save the connection profile and remote access VPN policy. Deploy the remote access VPN on Firewall Threat Defense.

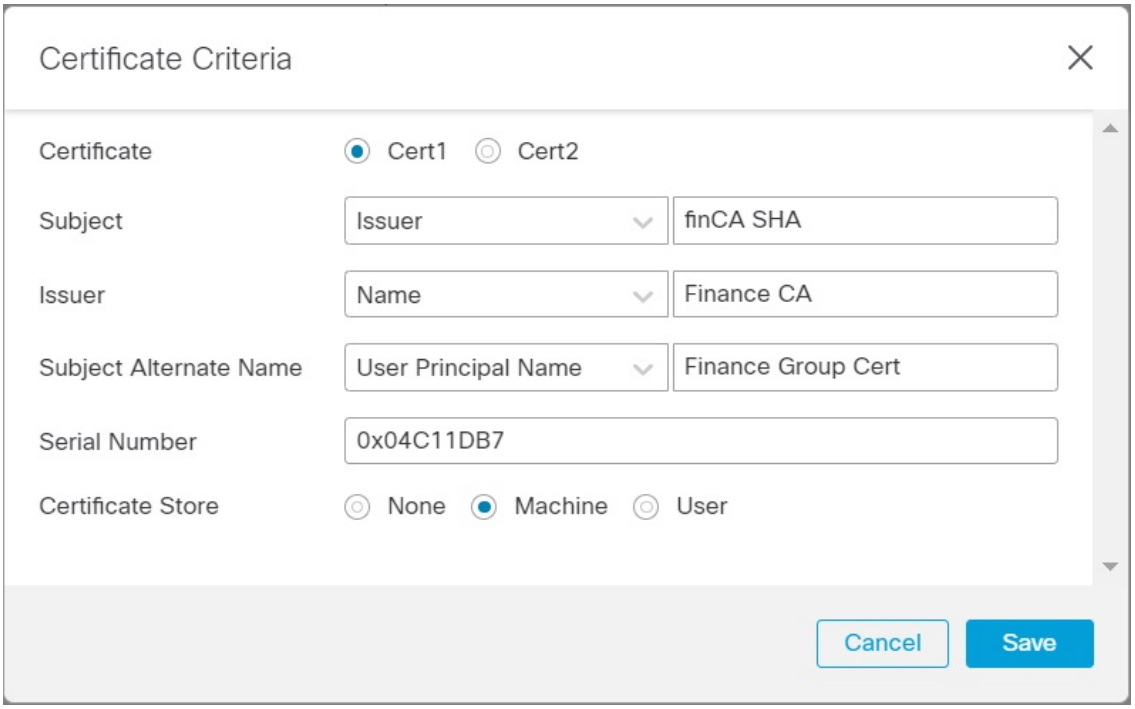
For information about remote access VPN AAA settings, see Configure AAA Settings for Remote Access VPN, on page 1835.

Certificate Configuration in DAP

You can also configure certificate criteria attributes in a DAP record. The user and machine certificate received from the VPN client during multiple-certificate authentication is loaded into dynamic access policy (DAP) to allow policies to be configured based on the field of the certificate. You can make policy decisions based on the fields of a certificate used to authenticate that connection attempt.

1. Choose **Devices** > **VPN** > **Dynamic Access Policy**.
2. Edit an existing DAP policy or create a new one and then edit the policy.
3. Choose an existing DAP record, or create a new one and then edit the record.
4. Select **Endpoint Criteria** > **Certificate**.
5. Select the Match Criteria **All** or **Any**.
6. Click **Add** to add certificate attributes.

Figure 377:



7. Select the certificate, **Cert1** or **Cert2**.
8. Select the **Subject** and specify the certificate subject value.

9. Select the **Issuer** and specify the certificate issuer name.
10. Select the **Subject Alternate Name** and specify the alternate name for the subject.
11. Specify the **Serial Number**.
12. Select the **Certificate Store**: None, Machine, or User.

    This option adds a condition to check for the store from which the certificate is picked on the endpoint.
13. Click **Save** to complete the certificate criteria settings.

    Configure the required DAP record settings and then associate the DAP with the remote access VPN.

For more information about DAP, see Dynamic Access Policies , on page 1927.

Monitor Remote Access VPNs

Remote Access VPN Dashboard

Remote Access Virtual Private Network (RA VPN) allows remote users to securely connect to your network. The RA VPN dashboard allows you to monitor real-time data from active RA VPN sessions on the devices. You can quickly determine problems related to user sessions and mitigate the problems for your network and users.

RA VPN dashboard (**Overview** > **Dashboards** > **Remote Access VPN**) provides a snapshot of the active RA VPN sessions on the threat defense devices managed by the Cloud-Delivered Firewall Management Center.

The dashboard provides:

- Visualization of active user sessions based on their location.
- Detailed information about the active user sessions.
- Mitigation of user session problems by terminating sessions, if required.
- Distribution of active user sessions per device, encryption type, Secure Client version, operating system, and connection profile.
- Device identity certificate expiration details of the devices.

The dashboard has the following widgets:

- Active Sessions (Tabular View)
- Active Sessions (Map view)
- Sessions
- Device Identity Certificates

Active Sessions (Tabular View)

This widget provides a tabular view of the active RA VPN users connected. You can view details of the active RA VPN sessions such as username, assigned IP, public IP, login time, VPN gateway (threat defense device), client application, client operating system, connection profile, and group policy. You can use the filter to

narrow down your search based on the different criteria. You can also perform the following actions on the individual sessions:

- Terminate a session of a specific user.
- Terminate all sessions of a specific user connected to a specific VPN gateway.
- Terminate all sessions that are connected to a specific VPN gateway.

If a client device supports dual address stack and the RA VPN configuration on the Firewall Threat Defense device allows IPv4 and IPv6 address pools, when a client establishes an RA VPN session with the headend device, it assigns an IPv4 and an IPv6 address to the client's tunnel interface. The RA VPN session has two IP addresses, an IPv4 and an IPv6 address on the threat defense device. The Cloud-Delivered Firewall Management Center shows two sessions for the same user, one with an IPv4 address and another with an IPv6 address, and the session count is two.

Hence, even when there is only a single RA VPN session from a user as per **show vpn-sessiondb l2l filter ipaddress** command on the device, the Cloud-Delivered Firewall Management Center shows two different sessions.

Active Sessions (Map View)

This widget shows an interactive heat map to visualize the location of the users connected through RA VPN sessions on the devices.

- Countries that have user sessions appear in shades of blue.
- Legend of the map provides a scale that indicates the correlation between the number of sessions in a country and the shade of blue for the country.
- Hover the mouse pointer over the map to view the country name and the total number of active user sessions.
- Zoom in, zoom out, and reset options are available.

Sessions

This widget allows you to monitor real-time data from active RA VPN sessions on the devices. You can filter and view the distribution of active RA VPN sessions according to:

- Device: Displays the number of sessions per device.
- Encryption Type: Displays the number of Secure Client SSL or IPsec sessions.
- Secure Client Version: Displays the sessions per Secure Client version.
- Operating System: Displays the sessions per operating systems. For example, Windows, Linux, Mac, Mobile OS, and so on.
- Connection Profile: Displays the sessions per connection profile.

Device Identity Certificates

This widget provides information about the identity certificate expiry of the RA VPN gateways. You can view expired certificates and certificates that are due for expiry within a month. Click **View Details** to view the certificates in the **Devices** > **Certificates** page.

VPN Session and User Information

The system generates events that communicate the details of user activity on your network, including VPN-related activity. The system monitoring capabilities enable you to determine quickly whether remote access VPN problems exist and where they exist. You can then apply this knowledge and use your network management tools to reduce or eliminate problems for your network and users. Optionally, you can log out remote access VPN users as needed.

Viewing Remote Access VPN Active Sessions

Analysis > **Users heading** > **Active Sessions**

Lets you view the currently logged-in VPN users at any given point in time with supporting information such as the user name, login duration, authentication type, assigned/public IP address, device details, client version, endpoint information, throughput, bandwidth consumed group policy, tunnel group and so on. The system allows you to filter current user information, log users out, and delete users from the summary list.

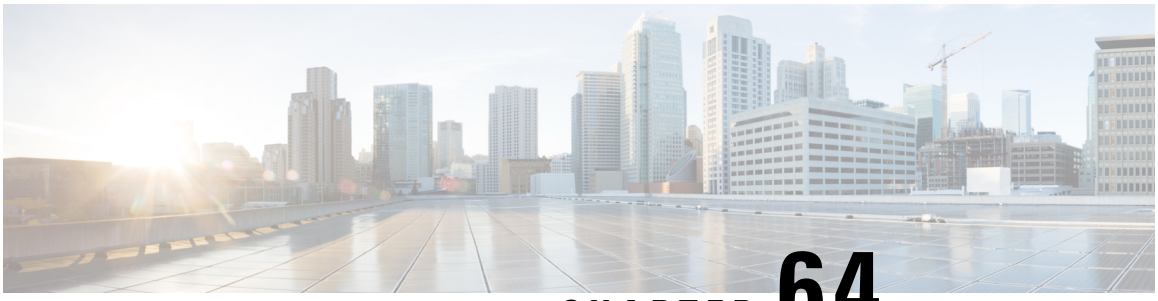


Note If you configure your VPN in a high-availability deployment, the device name displayed against active VPN sessions can be the primary or secondary device that identified the user session.

Viewing Remote Access VPN User Activity

Analysis > **Users heading** > **User Activity**

Lets you view the details of user activity on your network. The system logs historical events and includes VPN-related information such as connection profile information, IP address, geolocation information, connection duration, throughput, and device information.



CHAPTER 64

Dynamic Access Policies

Dynamic access policies (DAP) enable you to configure authorization that addresses the dynamics of VPN environments. You create a dynamic access policy by setting a collection of access control attributes that you associate with a specific user tunnel or session. These attributes address issues of multiple group membership and endpoint security.

- Firewall Threat Defense dynamic access policies, on page 1927
- Prerequisites for dynamic access policy, on page 1929
- Guidelines and limitations for dynamic access policies, on page 1929
- Configure a Dynamic Access Policy (DAP), on page 1930
- Associate dynamic access policy with remote access VPN, on page 1941

Firewall Threat Defense dynamic access policies

A Secure Firewall Threat Defense dynamic access policy is a collection of access control attributes that

- addresses issues of multiple group memberships and endpoint security in VPN environments
- grants access to particular users for particular sessions according to defined policies, and
- adapts to dynamic environments with multiple variables affecting each VPN connection.

Dynamic access policy operation

VPN gateways operate in dynamic environments. Multiple variables can affect each VPN connection. For example, intranet configurations that frequently change, the various roles each user inhabits within an organization, and log in attempts from remote access sites with different configurations and levels of security. The task of authorizing users is much more complicated in a VPN environment than it is in a network with a static configuration.

You can create a dynamic access policy by setting a collection of access control attributes that you associate with a specific user tunnel or session. The Firewall Threat Defense device generates a DAP during user authentication by selecting or aggregating attributes from one or more DAP records. The device then selects these DAP records based on the endpoint security information of the remote device and AAA authorization information for the authenticated user. Then the device applies the DAP record to the user tunnel or session.

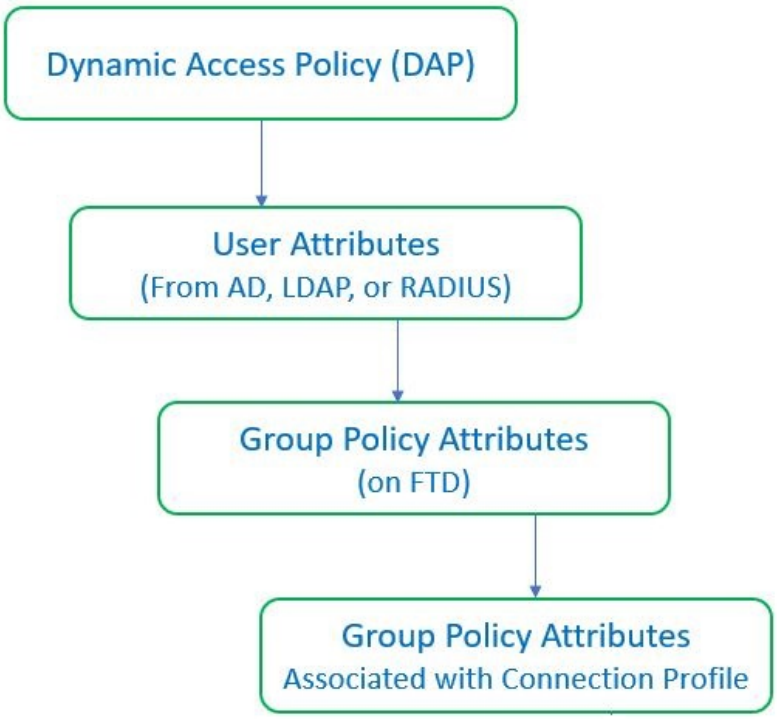
Hierarchy of policy enforcement of permissions and attributes in a Firewall Threat Defense device

The Firewall Threat Defense device supports applying user authorization attributes, also called user entitlements or permissions, to VPN connections. The attributes are applied from a DAP on the Firewall Threat Defense, external authentication server and/or authorization AAA server (RADIUS) or from a group policy on the Firewall Threat Defense device.

If the Firewall Threat Defense device receives attributes from all sources, the device evaluates, merges, and applies the attributes to the user policy. If there are conflicts between attributes coming from the DAP, the AAA server, or the group policy, the attributes from the DAP always take precedence.

The Firewall Threat Defense device applies attributes in this order:

Figure 378: Policy enforcement flow



1. **DAP attributes on the FTD**—The DAP attributes take precedence over all others.
2. **User attributes on the external AAA server**—The server returns these attributes after successful user authentication and/or authorization.
3. **Group policy configured on the FTD**—If a RADIUS server returns the value of the RADIUS Class attribute IETF-Class-25 (OU= group-policy) for the user, the Firewall Threat Defense device places the user in the group policy of the same name and enforces any attributes in the group policy that are not returned by the server.

4. **Group policy assigned by the Connection Profile (also known as Tunnel Group)**—The Connection Profile has the preliminary settings for the connection, and includes a default group policy that is applied to the user before authentication.



Note The Firewall Threat Defense device does not support inheriting system default attributes from the default group policy, *DfltGrpPolicy*. For the user session, the device uses the attributes on the group policy that you assign to the connection profile, unless the user attributes or the group policy from the AAA server overrides them.

Prerequisites for dynamic access policy

This reference provides the licensing prerequisites that must be met before configuring Dynamic Access Policy features.

Licensing prerequisites

- Firewall Threat Defense must have at least one of these Secure Client licenses:
  - Secure Client Premier
  - Secure Client Advantage
  - Secure Client VPN Only
- The Firewall Threat Defense Essentials license must allow export-controlled functionality.

General prerequisite

Ensure that you have the Secure Firewall Posture package before you configure the dynamic access policy. You can add the Secure Firewall Posture file at **Objects** > **Object Management** > **VPN** > **Secure Client File**.

Guidelines and limitations for dynamic access policies

Consider these guidelines and limitations when implementing Dynamic Access Policies:

- Matching of AAA attributes in a DAP will work only if a AAA server is configured to return the correct attributes when authenticating or authorizing a remote access VPN session.
- Minimum Secure Client and Secure Firewall Posture package version supported for DAP is 4.6. But it is highly recommended to use the latest version of Secure Client.
- DAP does not support clustering or multi-instance mode.
- DAP condition with an assigned IPv4 or IPv6 address does not work for local authentication.

Configure a Dynamic Access Policy (DAP)

Create a dynamic access policy

Create a dynamic access policy to enable conditional access control for VPN users based on their endpoint security posture and compliance status.

Dynamic access policies evaluate endpoint attributes and security posture to determine appropriate access permissions for VPN users. These policies work in conjunction with posture assessment packages to enforce security compliance.

Before you begin

Ensure that you have the Secure Firewall Posture package before you configure the dynamic access policy. You can add the Secure Firewall Posture file at **Objects** > **Object Management** > **VPN** > **Secure Client File**.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**, and click **Create Dynamic Access Policy**.

Step 2 Specify the **Name** for the DAP policy and an optional **Description**.

Step 3 Choose the **HostScan Package** from the drop-down list.

Step 4 Click **Save**.

What to do next

To configure DAP record, see Create Dynamic Access Policy Record

Create a dynamic access policy record

A dynamic access policy (DAP) can contain multiple DAP records, where you configure user and endpoint attributes. You can prioritize the DAP records within a DAP so that the Firewall Threat Defense device can select and sequence the required criteria when a user attempts VPN connection.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Edit an existing dynamic access policy or click **Create Dynamic Access Policy** to create a new one and then edit the policy.

Step 3 Click **Create DAP Record**.

Step 4 Click the **General** tab.

Step 5 Specify the **Name** for the DAP record.

Step 6 Enter the **Priority** for the DAP record.

The lower the number, the higher the priority.

Step 7 Select one of these actions to take when a DAP record matches:

- **Continue**—Applies access policy attributes to the session. If any, the next DAP record (next policy line with less priority) is then evaluated.
- **Terminate**—Terminates the session.
- **Quarantine**—Quarantines the connection.

Step 8 Check the **Display User Message on Criterion Match** check-box and add the user message.

The Firewall Threat Defense displays this message to the user when the DAP record matches.

Step 9 Check the **Apply a Network ACL on Traffic** check box and select the access control list from the drop-down.

Step 10 Check the **Apply one or more Secure Client Custom Attributes** check box and select the custom attributes object from the drop-down.

Step 11 Click **Save**.

Configure posture assessment criteria

For a DAP policy, you can configure file, process, or registry endpoint attributes with unique endpoint IDs. These IDs can be used as endpoint criteria in Lua scripts to configure a DAP record.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Click **Create Dynamic Access Policy** to create a new DAP policy. and then edit the policy.

Step 3 Click the edit icon adjacent to the DAP policy.

Step 4 Click **Add Posture Assessment Criteria**.

Step 5 Do one of these:

- Configure a file endpoint attribute:
  - **a.** Click the **File** radio button.
  - **b.** In the **Endpoint ID** field, enter a unique ID for the file. It can be a string or a number.
  - **c.** In the **File Path** field, specify the file path.
- Configure a registry endpoint attribute:
  - **a.** Click the **Registry** radio button.
  - **b.** In the **Endpoint ID** field, enter a unique ID for the registry. It can be a string or a number.
  - **c.** In the **Entry Path** field, specify the file path.

- Configure a process endpoint attribute:
  a. Click the **Process** radio button.
  b. In the **Endpoint ID** field, enter a unique ID for the process. It can be a string or a number.
  c. In the **Process Name** field, specify the process name.

Note
You cannot edit an Endpoint ID once it is saved.

Step 6 Click **Save**.

What to do next

You can use the endpoint IDs to configure advanced posture assessment criteria using Lua scripts. For more information, see Configure advanced settings for DAP, on page 1941.

Configure AAA criteria settings for a DAP record

Configure AAA criteria settings to specify the conditions under which DAP records are selected and applied to user sessions based on AAA authorization and posture assessment information.

DAP complements AAA services by providing a limited set of authorization attributes that can override the attributes that AAA provides. The Firewall Threat Defense device selects DAP records based on the AAA authorization information for the user and posture assessment information for the session. The Firewall Threat Defense device can choose multiple DAP records depending on this information, which it then aggregates to create DAP authorization attributes.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.
Step 2 Edit an existing DAP policy or create a new one and then edit the policy.
Step 3 Select a DAP record or create a new one, and edit the DAP record.
Step 4 Click **AAA Criteria**.
Step 5 Select one of the **Match criteria between sections**.

- Any—Matches any of the criteria.
- All—Matches all the criteria.
- None—Matches none of the set criteria.

Step 6 Click **Add** to add the required **Cisco VPN Criteria**.

Cisco VPN criteria include attributes for group policy, assigned IPv4 address, assigned IPv6 address, connection profile, username, username 2, and SCEP required.

a) Select an attribute and specify the **Value**.

b) Click **Add another criteria** to add more criteria.
c) Click **Save**.

Step 7 Select **LDAP Criteria**, **RADIUS Criteria**, or **SAML Criteria** and specify the **Attribute ID** and **Value**.

Step 8 Click **Save**.

Configure endpoint attribute selection criteria in DAP

Endpoint attributes contain information about the endpoint system environment, posture assessment results, and applications. A Firewall Threat Defense device dynamically generates a collection of endpoint attributes during session establishment and stores these attributes in a database that is associated with the session. Each DAP record specifies the endpoint selection attributes that must be satisfied for the Firewall Threat Defense device to choose it for a session. The Firewall Threat Defense device selects only DAP records that satisfy every condition configured.

Note When using remote Access VPN with DAP or Secure Firewall Posture (formerly HostScan) and change of authorization (CoA), endpoint attribute checks differ between initial authentication and CoA-triggered authorization.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**, and click **Create Dynamic Access Policy**.

Step 2 Create or edit a DAP record.

Step 3 Click **Endpoint Criteria** and configure attributes such as anti-malware, device, Secure Client, NAC, application, personal firewall, operating system, process, registry, file, and certificate authentication.

Note
You can create multiple instances of each type of endpoint attribute. There is no limit for the number of endpoint attributes for each DAP record.

Step 4 Click **Save**.

Add an anti-malware endpoint attribute to a DAP record

Add anti-malware endpoint attributes to a DAP to control access based on the presence and configuration of anti-malware software on client devices.

Anti-malware endpoint attributes allow you to create access policies that verify client devices have appropriate anti-malware protection before granting network access.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **Anti-Malware**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click **+** to add the anti-malware attributes.

In the **Anti-Malware** dialog box, configure these parameters:

a) Check the **Installed** check box to indicate whether the selected endpoint attribute and its accompanying qualifiers are installed or not installed.
b) In **Real Time Scanning**, choose **Enabled** or **Disabled** to activate or deactivate real-time malware scanning.
c) Choose the anti-malware vendor from the **Vendor** drop-down list.
d) Choose the anti-malware description from the **Product Description** drop-down list.
e) Choose the **Version** of the anti-malware product.
f) Specify the number of days since the **Last Update**.

You can indicate that an anti-malware update must occur in less than (<) or more than (>) the number of days you specify.

g) Click **Save**.

Add a device endpoint attribute to a DAP record

Add device endpoint attributes to a DAP to control access based on specific device characteristics such as hostname, MAC address, BIOS serial number, or port configuration.

Device endpoint attributes allow you to create policies that evaluate device-specific criteria to determine access permissions. These attributes help enforce security policies based on device identity and configuration.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **Device**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click **+** to add application attributes.

In the **Device** dialog box, configure these parameters:

Select the **=** or **≠** operator to check the attribute to be equal to or not equal to the value you enter for these attributes:

- **Host Name**—Hostname of the device. Use the computer's host name only, not the fully qualified domain name (FQDN).

- **MAC Address**—MAC address of the network interface card you are testing for. The address must be in the format xxxx.xxxx.xxxx where x is a hexadecimal character.
- **BIOS Serial Number**—BIOS serial number value of the device you are testing for. The number format is manufacturer-specific.
- **Port Number**—Listening port number of the device.
- **Secure Desktop Version**—Version of the Secure Firewall Posture image running on the endpoint.
- **OPSWAT Version**—The OPSWAT client version.
- **Privacy Protection**—None, Cache cleaner, Secure Desktop.
- **TCP/UDP Port Number**—TCP or UDP port in the listening state.

Step 6 Click **Save**.

Add Secure Client endpoint attributes to a DAP record

Add specific endpoint attributes to a Dynamic Access Policy (DAP) record to control client access based on Secure Client characteristics such as client version, platform, and device identifiers.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **Secure Client**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click + to add the attributes.

In the **Secure Client** dialog box, configure these parameters:

Select the = or ≠ operator to check the attribute to be equal to or not equal to the value you enter.

a) Specify the **Client Version**.
b) Choose the platform from the **Platform** drop-down list.
c) Choose the platform version from the **Platform Version** drop-down list.
d) Specify the **Device Type** and **Device Unique ID**.
e) Add the **MAC Addresses**.

Note
The MAC Address must be in the format XX-XX-XX-XX-XX-XX, where each X is a hexadecimal character. You can click **Add another MAC Address** to add more addresses.

f) Click **Save**.

Add NAC endpoint attributes to a DAP record

You can configure Network Access Control (NAC) endpoint attributes within a DAP (Dynamic Access Policy) record to evaluate endpoint posture status based on a defined criteria.

NAC endpoint attributes in DAP records help determine access permissions based on endpoint posture compliance. You can specify whether endpoints must match all criteria or any criteria, and define specific posture token strings for evaluation.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **NAC**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click **+** to add NAC attributes.

Step 6 In the **NAC** dialog box, set the operator to be equal to **=** or not equal to **≠** the posture status.

Step 7 Choose one of these options from the **Posture Status** drop-down list:

- Compliant
- Non-Compliant
- Unknown

Step 8 Click **Save**.

Add an application attribute to a DAP record

This task allows you to define application-specific criteria within a DAP record to control access based on the type of remote access client.

Application attributes in DAP records help determine access policies by evaluating the client application type used for remote connections. You can configure these attributes to match all or any specified criteria.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **Application**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click **+** to add application attributes.

Step 6 In the **Application** dialog box, choose equals (**=**) or does not equal (**≠**) and specify the **Client Type** to indicate the type of remote access connection. Choose from one of these options:

- Clientless

- Cut-Through-Proxy
- Secure Client
- IPsec
- L2TP
- IPsec-IKEv2-Generic-RA

Step 7 Click **Save**.

Add a personal firewall endpoint attribute to a DAP record

Adding personal firewall endpoint attributes to a DAP allows you to control network access based on whether client devices have specific firewall software installed, enabled, and meeting version requirements.

Personal firewall endpoint attributes help enforce security policies by ensuring that connecting devices meet firewall protection standards before granting network access.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **Personal Firewall**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click **+** to add personal firewall attributes.

In the **Personal Firewall** dialog box, configure these parameters:

a) Click **Installed** to indicate whether the personal firewall endpoint attribute and its accompanying qualifiers (fields below the Name/Operation/Value column) are installed or not installed.
b) Select **Enabled** or **Disabled** to activate or deactivate firewall protection.
c) Choose the name of the firewall vendor from the **Vendor** drop-down list.
d) Choose the firewall description from the **Product Description** drop-down list.
e) Select the equals (**=**) or does not equal (**≠**) operator and choose the **Version** of the personal firewall product.
f) Click **Save**.

Add an operating system endpoint attribute to a DAP record

Operating system endpoint attributes help create granular access policies by identifying and filtering client devices based on their operating system characteristics. This is useful for implementing different access levels or restrictions based on device types.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **Operating System**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click **+** to add endpoint attributes.

In the **Operating System** dialog box, configure these parameters:

a) Select the equals (**=**) or does not equal (**≠**) operator and then choose the operating system from the **Operating System** drop-down list.
b) Select the equals (**=**) or does not equal (**≠**) operator and then specify the operating system version in the **Version** field.
c) Click **Save**.

Add a process endpoint attribute to a DAP

Configure process endpoint attributes to define specific criteria for device access policy enforcement based on running processes.

Process endpoint attributes allow you to create device access policies based on specific processes running on endpoints. This helps ensure that only devices with approved processes can access network resources.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **Process**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click **+** to add process attributes.

In the **Process** dialog box, configure these parameters:

a) Select **Exists** or **Does not exist**.
b) Specify the **Process Name**.
c) From the **Endpoint ID** drop-down list, choose the process ID or click **+** to configure a posture assessment criteria for the process. For more information, see Configure posture assessment criteria, on page 1931.
d) For **Process Name**, click **Exists** or **Does not exist** and specify the process name.
e) Click **Save**.

Add a registry endpoint attribute to a DAP record

Registry endpoint attributes allow you to define specific registry keys that the system scans during posture assessment to determine endpoint compliance status.

Scanning for registry endpoint attributes applies to Windows operating systems only.

Before you begin

Before configuring a registry endpoint attribute, define the registry key for which you want to scan in the Secure Firewall Posture window for Cisco Secure Desktop.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **Registry**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click + to add registry attributes.

In the **Registry** dialog box, configure these parameters:

a) Select the **Entry Path** for the registry and specify the path.
b) From the **Endpoint ID** drop-down list, choose the ID for the registry or click + to configure a posture assessment criteria for the registry. For more information, see Configure posture assessment criteria, on page 1931.
c) Choose the existence of the registry, **Exists** or **Does not exist**.
d) Select the registry **Type** from the list.
e) Select the equals (=) or does not equal (≠) operator and enter the **Value** of the registry key.
f) Select **Case insensitive** to disregard the case of the registry entry while scanning.
g) Click **Save**.

Add a file endpoint attribute to a DAP record

You can configure file endpoint criteria within a DAP (Device Access Policy) record to specify file attributes used for posture assessment and device authorization.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP policy and a DAP record.

Step 3 Click the **Endpoint Criteria** tab and click **File**.

Step 4 Select the **Match criteria** as **All** or **Any**.

Step 5 Click + to add file attributes.

In the **File** dialog box, configure these parameters:

a) From the **Endpoint ID** drop-down list, choose the ID for the file or click **+** to configure a posture assessment criteria for the file. For more information, see Configure posture assessment criteria, on page 1931.
b) Specify the **File Path**.
c) Select **Exists** or **Does not exist** to indicate the presence of the file.
d) Select less than (<) or greater than (>) and specify the **Last Modified** days for the file.
e) Select the equal to (=) or not equal to ≠ operator and enter the **Checksum**.
f) Click **Save**.

Add certificate authentication attributes to a DAP record

Adding certificate authentication attributes to a DAP allows you to control access based on specific certificate fields such as subject, issuer, and serial number.

You can index each certificate to allow referencing to any of the received certificates, by the configured rules. Based on these certificate fields, you can configure DAP rules to allow or disallow connection attempts.

Note This configuration is applicable only for certificates used with multiple certificate authentication. It does not apply for single client certificate authentication.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.
Step 2 Create or edit a DAP record.
Step 3 Click the **Endpoint Criteria** tab and click **Multiple Certificate Authentication**.
Step 4 Select the **Match criteria** as **All** or **Any**.
Step 5 Click **+** to add certification authentication criteria.

In the **Multiple Certificate Authentication Criteria** dialog box, configure these parameters:

a) Select **Cert1** or **Cert2** certificate.
b) Select the **Subject** and specify the subject value.
c) Select the **Issuer** and specify the issuer value.
d) Select the **Subject Alternate Name** and specify the subject value.
e) Specify the **Serial Number**.
f) Choose the **Certificate Store** as **None**, **Machine**, or **User**.

The VPN client sends the certificate store information.

g) Click **Save**.

Configure advanced settings for DAP

This task allows you to configure selection criteria for DAP records using Lua scripts when the standard AAA AND endpoint attribute areas are insufficient for your requirements.

You can use the **Advanced** tab for adding selection criteria other than what is possible in the AAA AND endpoint attribute areas. For example, while you can configure the Firewall Threat Defense to use AAA attributes that satisfy any, all, OR none of the specified criteria, the endpoint attributes are cumulative, AND must satisfy all. To let the security appliance employ one endpoint attribute OR another, you must create appropriate logical expressions in Lua AND enter them here.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Create or edit a DAP record.

Step 3 Click the **Advanced** tab.

Step 4 Select **AND** OR **OR** as the match criteria to use in the DAP configuration.

Step 5 Add the Lua script in the **Lua script for advanced attribute matching** field.

Step 6 To use the endpoint criteria ID in your Lua script:

a. Place the cursor at the point where you want to insert the endpoint criteria ID.

b. From the **Endpoint Criteria** drop-down list, choose the criteria a

c. Choose the corresponding ID from the adjacent drop-down list.

Example:

In this example, DAPTESTFILE, LIBAGENT, vpnagent, AND DUOAGENT were inserted in the Lua script:

```
EVAL(endpoint.file["DAPTESTFILE"].exists,"EQ","true") or
EVAL(endpoint.file["LIBAGENT"].exists,"EQ","true") and
EVAL(endpoint.process[""vpnagent""].exists,"EQ","true") and
EVAL(endpoint.registry[""DUOAGENT""].exists,"EQ","true")
```

Step 7 Click **Save**.

Associate dynamic access policy with remote access VPN

You can associate Dynamic Access Policy (DAP) with remote access VPN policy for the dynamic access policy attributes to match during VPN session authentication and authorization. You can then deploy the remote access VPN on the Firewall Threat Defense.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** next to the remote access VPN policy to which you want to associate dynamic access policy.

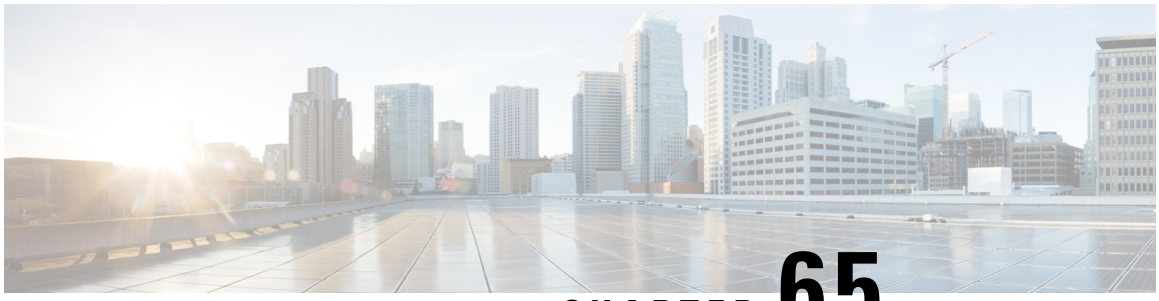
Step 3 Click the link in remote access VPN to select the dynamic access policy.

Step 4 Select the policy from the **Dynamic Access Policy** drop-down or click **Create a new Dynamic Access Policy** to configure a new dynamic access policy.

Step 5 Click **OK**.

Step 6 Click **Save** to save the remote access VPN policy.

When the remote access VPN user tries to connect, the VPN checks the configured dynamic access policy records and attributes. VPN creates a dynamic access policy based on the matching dynamic access policy records and takes appropriate action on the VPN session.



CHAPTER 65

VPN Monitoring and Troubleshooting in Security Cloud Control

- Monitor site-to-site VPNs using Site-to-Site VPN summary page, on page 1943
- Monitor Remote Access VPN Sessions, on page 1943
- SD-WAN summary dashboard, on page 1943

Monitor site-to-site VPNs using Site-to-Site VPN summary page

You can view a summary of your site-to-site VPN topologies in the **Site-to-Site VPN Summary** page. For each topology, you can view details such as VPN type, network topology, VPN interfaces, VPN devices, and tunnel statuses.

You can edit or delete the topology using the edit and delete buttons. For SASE topology VPNs, you have options to deploy, edit and delete any topology.

Monitor Remote Access VPN Sessions

The Security Cloud Control Remote Access Monitoring dashboard can be used to view consolidated information about RA VPN users, including the current status of users, device types, client applications, user geolocation information, and duration of connections. You can also disconnect RA VPN sessions as needed.

See Monitor Remote Access Virtual Private Network Sessions for more information.

SD-WAN summary dashboard

The SD-WAN Summary dashboard (**Overview** > **Dashboards** > **SD-WAN SummaryAnalysis** > **SD-WAN Summary**) provides a snapshot of your WAN devices and their interfaces. This dashboard helps you to:

- Identify issues with the underlay and overlay (VPN) topologies.
- Troubleshoot VPN issues using the existing **Health Monitoring**, **Device Management**, and **Site-to-Site Monitoring** pages.
- Monitor application performance metrics of WAN interfaces. The threat defense steers application traffic based on these metrics.

For clusters, this dashboard displays application performance metrics of only the control node and not the data nodes.

WAN device criteria

A WAN device must meet one of these criteria:

- The device must be a VPN peer.
- The device must have WAN interface.

WAN interface criteria

A WAN interface must meet one of these criteria:

- The interface has IP address-based path monitoring enabled on it.
- The interface has a Policy Based Routing (PBR) policy with at least one application configured to monitor it.

For more information about PBR policy and path monitoring, see Policy-Based Routing, on page 1541.

Uplink decisions

Click **Uplink Decisions** to view the **VPN Troubleshooting** page. You can view syslogs with ID: 880001. These syslogs show the threat defense interfaces through which it steers traffic based on the configured PBR policy.

To view the above syslogs and to view the data on this dashboard, ensure that you review Prerequisites for using SD-WAN summary dashboard, on page 1806.

Prerequisites for using SD-WAN summary dashboard

Ensure that you review these prerequisites before using the SD-WAN summary dashboard.

General prerequisites

- You must be an Admin, Security Analyst, or Maintenance user to view this dashboard. See Secure Firewall Management Center and Cloud-delivered Firewall Management Center User Role Mapping for more information.
- Threat Defense devices must be Version 7.2 or later.
- Enable IP-based path monitoring and HTTP-based application monitoring on the WAN interfaces.

  To do this:

  1. Choose **Devices** > **Device Management**.
  2. Click the edit icon adjacent to the device that you want to edit.
  3. Click the edit icon adjacent to the interface that you want to edit.
  4. Click the **Path Monitoring** tab.
  5. Check the **Enable IP based Monitoring** check box.

6. Check the **Enable HTTP based Application Monitoring** check box.
7. Click **OK**.

- Configure a PBR policy with at least one application configured to monitor it:
  1. Choose **Devices** > **Device Management**.
  2. Click the edit icon adjacent to the device that you want to edit.
  3. Click **Routing**.
  4. In the left pane, click **Policy Based Routing**.
  5. Click **Add**.
  6. From the **Ingress Interface** drop-down list, choose an interface.
  7. Click **Add** to configure a forwarding action.
  8. Configure the parameters.
  9. Click **Save**.
- To view syslogs when you click **Uplink Decisions**, you must:
  - Choose **Devices** > **Platform Settings** and create or edit a threat defense policy.
  - In the left pane, click **Syslog**.
  - Click the **Logging Setup** tab.
  - Check the **Enable Logging** check box to turn on the data plane system logging for the threat defense device.
  - Click the **All Logs** radio button to enable logging of all the troubleshooting syslog messages.

    or

    Click the **VPN Logs** radio button to enable logging of only the VPN troubleshooting messages.
  - Click **Save**.

Monitor WAN devices and interfaces using the SD-WAN summary dashboard

The SD-WAN Summary dashboard has these widgets :

- WAN connectivity, on page 1946
- VPN topology, on page 1946
- interface throughput, on page 1946
- Device inventory, on page 1946
- WAN device health, on page 1946

WAN connectivity

This widget provides a summary of the WAN interfaces statuses. It shows the number of WAN interfaces that are in the **Online**, **Offline** or **No Data** states. Note that you cannot monitor subinterfaces using this widget.

Click **View All Interfaces** to view more details about the interfaces in the health monitor page.

If a WAN interface is in the **Offline** or **No Data** state, you can troubleshoot it from the health monitor page:

1. In the **Monitoring** pane, expand **Devices**.
2. Click the corresponding WAN device to view the device-specific health details.
3. Click the **Interface** tab to view the interface status and aggregate traffic statistics for a specific time.

   Alternatively, you can click **View System & Troubleshoot Details**. The health monitor page is displayed with all the necessary details.

VPN topology

This widget provides a summary of the site-to-site VPN tunnel statuses. It shows the number of **Active**, **Inactive**, and **No Active Data** VPN tunnels.

Click **View All Connections** to view the VPN tunnel details in the **Site-to-site VPN Monitoring** dashboard.

If the tunnels are in the **Inactive** or **No Active Data** state, you can troubleshoot using the **Site-to-site VPN Monitoring** dashboard. In the **Tunnel Status** widget, hover your cursor over a topology, click **View** ( ) and do one of these actions:

- Click the **CLI Details** tab to view the details of the VPN tunnels.
- Click the **Packet Tracer** tab to use the packet tracer tool for the topology.

interface throughput

This widget monitors the average throughput of the WAN interfaces during the chosen time period.

The interface throughput is classified into four bands. These details aid in cost planning and resourcing. You can choose a time range for the widget data from the **Show Last** drop-down list. The range is from 15 minutes to two weeks.

Click **View Health Monitoring** to view more details about the interface in the health monitor page.

Device inventory

This widget lists all the managed WAN devices and groups them according to the model.

Click **View Device Management** to view more details about the device in the **Device Management** page.

WAN device health

This widget displays the device count according to the health of the WAN devices. You can view the number of devices with errors, warnings, or those that are in **Disabled** state.

Click **View Health Monitoring** to view the alarms, and quickly identify, isolate, and resolve issues.

If the health of a device is affected, you can troubleshoot it from the health monitor page.

1. In the **Monitoring** pane, expand **Devices**.

2. Click the corresponding WAN device to view the device-specific health details.
3. Click **View System & Troubleshoot Details**. The health monitor page is displayed with all the necessary details.

A device can be in **Disabled** state for multiple reasons, including these:

- Management interface is disabled.
- Device is powered off.
- Device is being upgraded.

Monitor application performance metrics of WAN interfaces using SD-WAN summary dashboard

You can monitor application performance metrics of WAN interfaces using the **Application Monitoring** tab of the SD-WAN Summary dashboard. These metrics include Jitter, Round Trip Time (RTT), Mean Opinion Score (MOS), and Packet Loss.

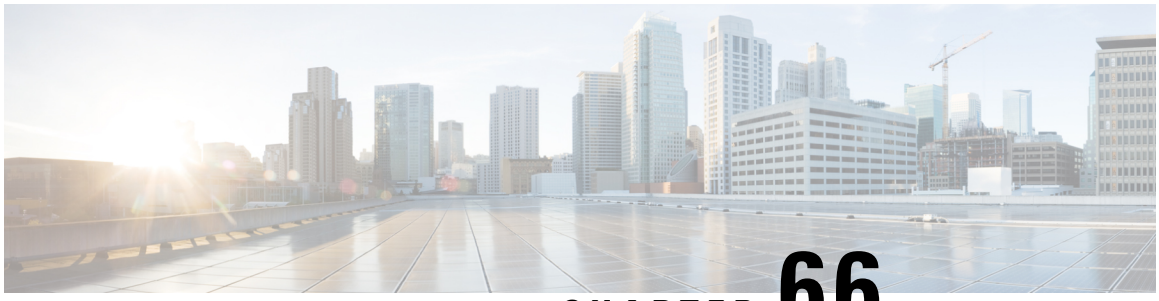
By default, the metrics data is refreshed every 5 minutes. You can change the refresh time; the range is from 5 to 30 minutes. You can view the metrics in tabular and graphical formats. For each WAN interface, the latest metric value appears in the table. For graphical data, you can choose a time interval of up to 24 hours to view the metrics data for the corresponding WAN interfaces.



PART XVI

Zero Trust Network Access

- Zero Trust Network Access Overview, on page 1951
- Clientless Zero Trust Network Access, on page 1955
- Universal Zero Trust Network Access, on page 1975



CHAPTER 66

Zero Trust Network Access Overview

- Clientless and universal zero trust network access, on page 1951
- Zero trust network access, on page 1952
- Comparison of clientless and universal ZTNA, on page 1952
- History for zero trust network access, on page 1953

Clientless and universal zero trust network access

Zero Trust Access or Zero Trust Network Access (ZTNA) is a security process that protects internal resources using identity-based access, including user trust and posture. The process verifies the user, device, and request context before allowing access to internal resources and applications. It grants the least privilege access only after verifying the user, understanding the request context, and assessing any risks.

ZTNA deployment models

You can choose between two ZTNA deployment models.

- Clientless Zero Trust Network Access
- Universal Zero Trust Network Access

Clientless zero trust network access

Clientless Zero Trust Network Access (Clientless ZTNA) authenticates and authorizes access to protected web-based resources and applications from inside (on-premises) or outside (remote) the network. Clientless ZTNA accomplishes this authentication using an external identity provider (IdP) with Security Assertion Markup Language (SAML) authentication. No client software is needed on the user's device.

Clientless ZTNA is ideal for remote users accessing web applications and unmanaged devices.

Universal zero trust network access

Universal Zero Trust Network Access (Universal ZTNA) is a client-based solution that provides identity-based access to all internal resources regardless of user location. It enforces strong authentication, posture validation, and traffic inspection for each application and user.

Universal ZTNA supports both remote and on-premises users. On-premises users access a trusted network using a Secure Firewall Threat Defense device. Remote users access a trusted network through the cloud-based Secure Access service, which evaluates policies and proxies the user traffic securely in the cloud.

Zero trust network access

Zero Trust Network Access (ZTNA) provides secure, segmented application access by requiring explicit authorization for every connection. This approach reduces the attack surface and simplifies secure migrations without adding devices.

Key characteristics

ZTNA uses your existing Threat Defense deployment as an enforcement point for application access, enabling segmented and per-application authorization. ZTNA establishes individual tunnels for both remote and on-premises users.

ZTNA hides the network from users and ensures that they can only access applications for which they have been explicitly authorized. Authorization for one application does not provide implicit access to other applications, which significantly reduces the attack surface. In other words, every access to an application or a resource must be explicitly authorized.

You can adopt a more secure access model by adding zero trust access functionality to Threat Defense. This eliminates the need to install or manage additional devices in the network.

Comparison of clientless and universal ZTNA

This comparison outlines the key technical and operational differences between clientless ZTNA and universal ZTNA deployment models. Clientless ZTNA uses a browser-based approach, which is ideal for guest users or unmanaged devices that require access to web applications. Universal ZTNA uses an agent-based solution that supports all protocols. It continuously monitors device posture in managed corporate environments. Consider factors such as application support, security inspection capabilities, and licensing requirements to choose the best strategy for securing your remote and on-premises workforce.

Factor	Clientless ZTNA	Universal ZTNA
Client required	No client is required(browser-based)	Yes, Secure Client is required
Application support	Supports web applications only	Supports all protocols and applications
Device posture	Device posture monitoring is limited (available through agents such as Duo Desktop)	Provides comprehensive and continuous device posture monitoring
Best for	Guest users, contractors, unmanaged devices	Employees, managed devices
Components involved	Third-party identity provider (IdP) configured with application access using Security Assertion Markup Language (SAML) authentication	• Secure Access • Secure Client software • Identity provider (depending on users' location)
Security policies	Deep inspection (Snort, IPS, File)	Deep inspection (Snort, IPS, File)

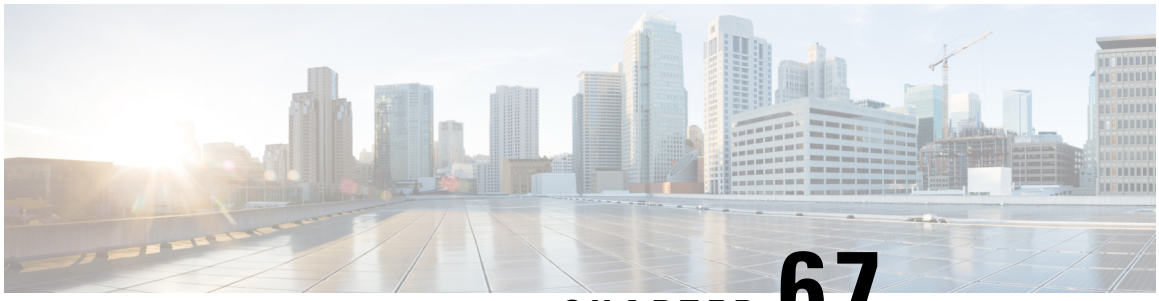
Factor	Clientless ZTNA	Universal ZTNA
User location	Remote and on-premises	Remote and on-premises
Licenses	Export-controlled license for Firewall Management Center	Licenses for Secure Access, Secure Client software, and Security Cloud Control Firewall Management (formerly Cisco Defense Orchestrator)

History for zero trust network access

This reference provides the minimum Firewall Management Center and Threat Defense versions needed to support zero trust application policies and monitoring.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Universal Zero Trust Network Access (universal ZTNA).	7.7.10	7.7.10	Universal Zero Trust Network Access (universal ZTNA) is a comprehensive solution that provides secure access to internal network resources based on user identity, trust, and posture. It ensures that access to one application does not implicitly grant access to the entire network, as with remote access VPN. New/modified screens: **Policies** > **Zero Trust Application** Requires Cisco Secure Access and Security Cloud Control. Deployment restrictions: Not supported with clustered devices, container instances, or transparent mode. Supported platforms: Secure Firewall 1150, 3100, 4100, 4200, and Firewall Threat Defense Virtual.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Clientless zero-trust access.	20230929	7.4.0 with Snort 3	Zero Trust Access allows you to authenticate and authorize access to protected web based resources, applications, or data from inside (on-premises) or outside (remote) the network using an external SAML Identity Provider (IdP) policy. The configuration consists of a Zero Trust Application Policy (ZTAP), Application Group, and Applications. New/modified screens: **Policies** > **Zero Trust Application** New/modified CLI commands: • **show running-config zero-trust application** • **show running-config zero-trust application-group** • **show zero-trust sessions** • **show zero-trust statistics** • **show cluster zero-trust statistics** • **clear zero-trust sessions application** • **clear zero-trust sessions user** • **clear zero-trust statistics**



CHAPTER 67

Clientless Zero Trust Network Access

Clientless Zero Trust Network Access enables users to access protected TLS web applications through their browser using SAML identity providers, eliminating the need to install a client on the endpoint. This solution leverages Zero Trust Application Policies, Application Groups, and applications. This document describes how to configure these features, outlines their limitations, and provides information on deployment, monitoring, and troubleshooting.

- Zero Trust Access, on page 1955
- Components of a Clientless Zero Trust Network Access configuration, on page 1956
- How Clientless Zero Trust Network Access works, on page 1957
- How threat defense works with zero trust access, on page 1958
- Prerequisites for zero trust application policy, on page 1960
- Limitations for Clientless Zero Trust Network Access, on page 1960
- Create a zero trust application policy, on page 1961
- Create an application group, on page 1963
- Create an application, on page 1965
- Create an ungrouped application, on page 1967
- Set targeted devices for zero trust access policy, on page 1968
- Edit a zero trust application policy, on page 1969
- Manage zero trust application policies, on page 1970
- Monitor zero trust sessions, on page 1971

Zero Trust Access

Clientless Zero Trust Network Access is based on Zero Trust Access principles. Zero Trust Access is a zero trust security model that eliminates implicit trust. The model grants the least privilege access after verifying the user, the context of the request, and after analyzing the risk if access is granted.

Clientless Zero Trust Network Access enables you to authenticate and authorize access to protected web-based resources and applications from inside (on-premises) or outside (remote) the network using an external Security Assertion Markup Language (SAML) identity provider (IdP) policy.

Key features

The solution includes these capabilities.

- Supports multiple SAML-based identity providers such as Duo, Azure Active Directory (Azure AD), Okta, and other identity providers.
- Client applications such as Cisco Secure Client are not required on the endpoint (client devices) for secure access.
- Access and authentication are performed through the browser.
- Supports only TLS web applications.
- Agents such as Duo Health support client device posture. These agents evaluate the posture against a policy in Duo and provide access based on this evaluation. Third-party identity providers such as Okta or PingID perform this same functionality with their agents that support posture evaluation.
- Supports HTTP-Redirect SAML binding.
- Supports application groups that make it easier to enable protection on a set of applications.
- Leverages threat defense intrusion and malware protection on zero trust application traffic.

You can use the Secure Firewall Management Center web interface to create a Zero Trust Application Policy that allows you to define private applications and assign threat policies to them. The policy is application specific where the administrator decides the inspection levels based on the threat perception for that application.

Components of a Clientless Zero Trust Network Access configuration

A new configuration of Clientless Zero Trust Network Access consists of a Zero Trust Application Policy, Application Group, and Applications.

Configuration components

The configuration includes these main components:

- **Zero Trust Application Policy**— consists of application groups, and grouped or ungrouped applications. Settings for Security Zones and Security Controls associate at a global level with all ungrouped applications and application groups.

  A global port pool is assigned to the policy, by default. A unique port is automatically assigned from this pool to each private application that is configured.

  Zero Trust Application policy consists of application groups, and grouped or ungrouped applications.

- **Application Groups**—Consists of a logical group of applications that share SAML authentication settings and can optionally share Security Zones and Security Controls settings.

  Application Groups inherit the Security Zones and Security Controls settings from the global policy and can override the values.

  When you create an Application Group, you can use the same SAML IDP configuration to authenticate multiple applications. Applications that are part of an Application Group inherit the Application Group's SAML configuration. This eliminates the need to configure the SAML settings for each application. After the Application Group is created, new applications can be added to it without configuring the IDP for it.

When an end user tries to access an Application that is part of group, the user is authenticated to the Application Group for the first time. When the user tries to access other applications that are part of the same Application Group, the user is provided access without being redirected again to the IDP for authentication. This prevents overloading the IDP with requests for application access and optimizes the usage of the IDP if a limit is enabled.

- **Applications**—There are two types:
  - **Ungrouped Applications**— Are standalone applications. SAML settings must be configured for every application. The applications inherit the Security Zones and Security Controls settings from the global policy and can be overridden by the application.
  - **Grouped Applications**— Are multiple applications that are grouped under an Application Group. The SAML settings are inherited from the Application Group and cannot be overridden. However, the Security Zones and Security Controls settings can be overridden for each application.

These certificates are required for the configuration.

- **Identity Certificate**—This certificate is used by Firewall Threat Defense to masquerade as the applications.Firewall Threat Defense behaves as a SAML Service Provider (SP). This certificate must be a wildcard or Subject Alternative Name (SAN) certificate that matches the FQDN of the private applications. It is a common certificate for all applications protected by Firewall Threat Defense.
- **IDP Certificate**—The IDP provides a certificate for each defined Application or Application Group. This certificate must be configured so that Firewall Threat Defense can verify the IDP's signature on incoming SAML assertions.



Note IDP certificates are commonly included within the metadata file; otherwise, users are required to have the IDP certificate readily available during the configuration of applications.

- **Application Certificate**—The encrypted traffic from user to the application is decrypted by Firewall Threat Defense using this certificate for the purpose of inspection.



Note This certificate is required to verify the cookies in the header to authorize connections, even if you are not conducting an IPS or Malware inspection.

How Clientless Zero Trust Network Access works

Summary

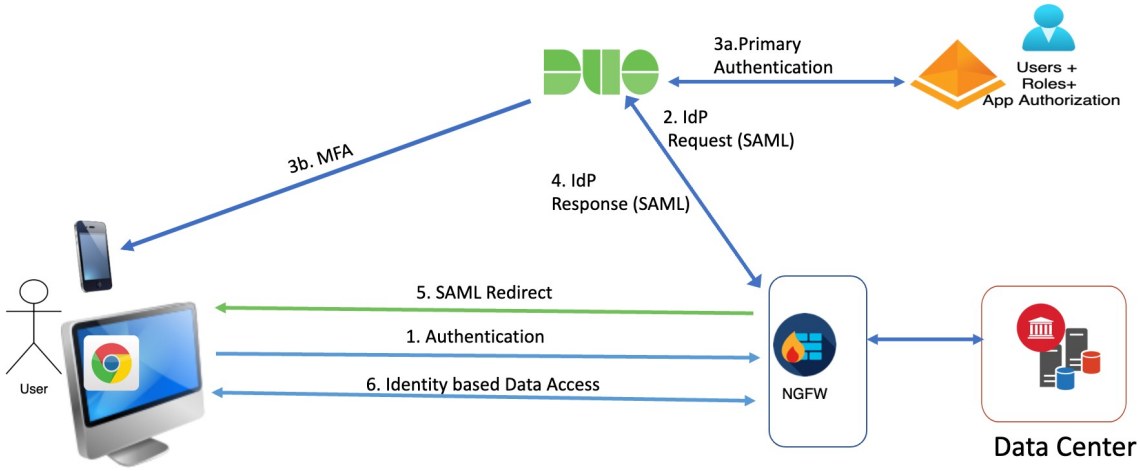
The key components involved in the Zero Trust Access clientless workflow are:

- User: Initiates access requests through a browser and provides authentication credentials.
- Identity Provider (IdP): Handles user authentication including primary and multi-factor authentication challenges.

- Threat Defense: Validates HTTPS requests, processes SAML responses, and controls application access.
- Application: The target resource that users access after successful authentication and authorization.

Workflow

Figure 379: Clientless zero trust network access workflow



These stages describe how the clientless zero trust network access workflow processes user requests:

1. User types the application URL in the browser and threat defense validates the HTTPS request.
   - If the HTTPS request is valid, the user is redirected to the mapped port.
   - If the HTTPS request is invalid, the user is sent for authentication per application.
2. The user is redirected to the configured identity provider (IdP).
3. The IdP performs authentication challenges for the user.
   - The user is redirected to the configured primary authentication source.
   - The user is challenged with the configured secondary multi-factor authentication, if any.
4. The IdP sends a SAML response to threat defense. The user ID and other necessary parameters are retrieved from the SAML response through the browser.
5. The user is redirected to the application.
6. The user is allowed access to the application after validation is successful.

How threat defense works with zero trust access

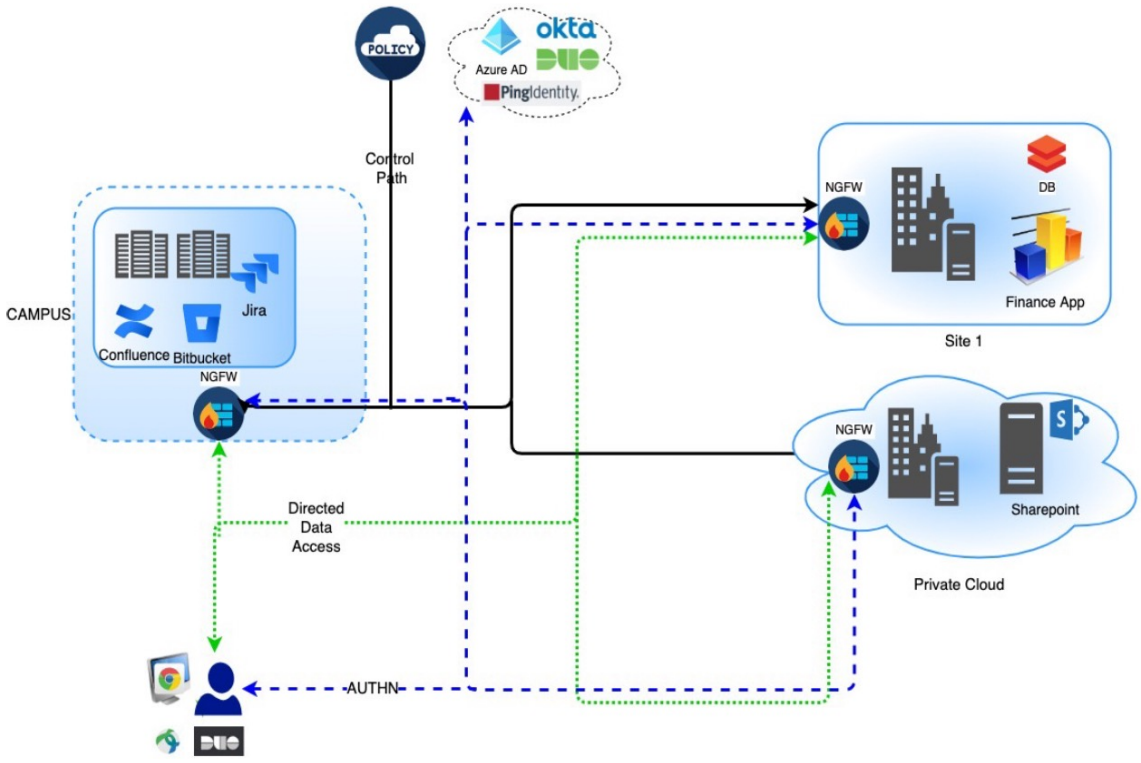
Summary

The key components involved in the Threat Defense Zero Trust Access process are:

- Remote or on-premises user: Initiates HTTPS requests from an endpoint using a browser to connect to applications.
- Threat Defense firewall: Intercepts user requests and redirects them to the identity provider for authentication and authorization.
- Identity Provider (IdP): Handles the authentication and authorization process for users attempting to access protected applications.
- Protected web applications: Applications secured behind the firewall that users can access after successful authentication.

Workflow

Figure 380: Threat defense deployment



These stages describe how Threat Defense works with Zero Trust Access:

1. Using a browser, a remote or on-prem user sends a HTTPS request to connect to an application from an endpoint.
2. The HTTPS request is intercepted by the firewall that protects the application.
3. The firewall redirects the user to application's configured IdP for authentication.

Note In the figure, each firewall protects a set of web applications. The user can directly access the applications behind the firewall after authentication and authorization.

4. After the authentication and authorization process is complete, the firewall allows the user to access the application.

Prerequisites for zero trust application policy

Licensing requirements

Ensure you have the necessary licensing before configuring Zero Trust Application Policy.

- Smart license account with export-controlled features
- (Optional) IPS and Threat licenses—You need these licenses if you use security controls.

Configuration requirements

Complete the required configurations for Zero Trust Application Policy implementation.

- Create a wildcard or Subject Alternative Name (SAN) certificate that matches the FQDN of private applications. For more information, see Add certificate enrollment objects, on page 2999.
- Create a security zone through which you can regulate access to private applications. For more information, see Create security zone and interface group objects, on page 1002.

Limitations for Clientless Zero Trust Network Access

Be aware of the following limitations when using Clientless Zero Trust Network Access:

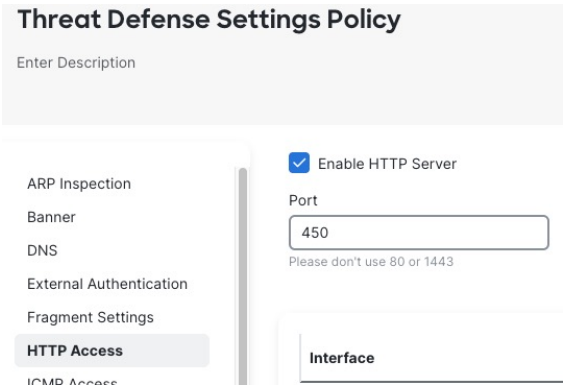
- Supports only TLS-enabled web applications. All traffic must be decrypted.
- Supports only SAML identity providers.
- IPv6 is not supported.

  NAT66, NAT64, and NAT46 scenarios are not supported.
- Supports IPv6; however, you can use only homogeneous network address translations (NAT) scenarios such as NAT66 and NAT44.

  NAT64, and NAT46 scenarios are not supported.
- Clientless zero trust network access is available on Firewall Threat Defense only if Snort 3 is enabled.
- By default, the Firewall Threat Defense device uses port 443 for secure communications. Because this is the same port used by TLS-enabled applications likely to be configured for zero trust network access, you must change the device's HTTP server port.

Go to **Devices** > **Platform Settings**, then edit the Threat Defense Settings policy for the device, click **HTTP Access**, select the **Enable HTTP Server** check box, and enter a port other than 443 in the field. When you're finished, click **Save** in the upper right corner of the page.

The following figure shows an example.



- All hyperlinks in protected web applications must have a relative path and are not supported on individual mode clusters.
- Protected web applications running on a virtual host or behind internal load balancers must use the same external and internal URL.
- Not supported on applications with strict HTTP Host Header validation enabled.
- If the application server hosts multiple applications and serves content based on the Server Name Indication (SNI) header in the TLS Client Hello, the external URL of the zero trust application configuration must match the SNI of that specific application.

Create a zero trust application policy

This task creates a Zero Trust Application Policy that enables secure access to private applications through authentication and authorization controls.

Before you begin

Ensure that you complete all the prerequisites listed in Prerequisites for zero trust application policy, on page 1960.

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**.

Step 2 Click **Add Policy**.

Step 3 In the **General** section, enter the policy name in the **Name** field. The description field is optional.

Step 4 Enter a domain name in the **Domain Name** field.

Ensure that the domain name is added to the DNS. The domain name resolves to the Firewall Threat Defense gateway interface from where the application is accessed. The domain name is used to generate the ACS URL for all private applications in an Application Group.

If you select an ACME certificate as the policy's identity certificate in the next step, the domain name must match the common name (CN) of the ACME certificate.

Note
When you change the domain name, the SAML Service Provider (SP) metadata gets updated. You must reconfigure these settings:

- IdP with the new SAML SP metadata
- DNS server with new domain name

When a deployment takes place after the domain name change, all the applications will be removed and readded, interrupting application access.

Step 5 Choose an existing certificate from the **Identity Certificate** drop-down list.

Click the **Add** (+) icon to configure a certificate enrollment object. For more information, see Add certificate enrollment objects, on page 2999.

You can choose an ACME certificate for authenticating the Firewall Threat Defense device as a SAML SP for a Zero Trust Application policy. ACME certificates automate the lifecycle management of SSL and TLS certificates, including their auto-renewal.

Step 6 Choose a security zone from the **Security Zones** drop-down list.

Click the **Add** (+) icon to add a new security zone.

To add security zones, see Create security zone and interface group objects, on page 1002.

Step 7 The **Global Port Pool** section sisplays a default port range. Modify it, if required. Port values range from 1024 to 65535. A unique port from this pool is assigned to each private application.

Note
- This port range should not conflict with the existing NAT range.
- Ensure that any intermediary or third-party firewall devices in your network are configured to allow traffic across the entire range defined in the **Global Port Pool**. Blocking these ports prevents application access after SAML authentication and authorization.

Step 8 (Optional) In the **Security Controls** section, you can add an Intrusion or Malware and File policy:

- **Intrusion Policy**—Choose a default policy from the drop-down list or click the **Add** (+) icon to create a new custom intrusion policy. For more information, see Creating a Custom Snort 3 Intrusion Policy topic in the latest version of the Cisco Secure Firewall Management Center Snort 3 Configuration Guide.
- **Variable Set**—Choose a default variable set from the drop-down list or click the **Add** (+) icon to create a new variable set. For more information, see Create a variable set, on page 3053.

  Note
  To use variable sets, you must have the Secure Firewall Threat Defense IPS license for your managed devices.

- **Malware and File Policy**—Choose an existing policy from the drop-down list. Click the **Add** (+) icon to create a new malware and file policy. For more information, see Manage file policies, on page 2601.

Step 9 Click **Save** to save the policy.

What to do next

1. Create an Application Group. See Create an application group, on page 1963.
2. Create an Application. See Create an application, on page 1965.
3. Associate a Zero Trust Application Policy with a device. See Set targeted devices for zero trust access policy, on page 1968
4. Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Create an application group

Create an application group to enable SAML-based authentication and manage zero trust network access applications with centralized security controls.

Application groups allow you to organize and manage multiple applications under a single SAML configuration, providing centralized authentication and security policy enforcement for zero trust network access.

Before you begin

Create a zero trust application policy, on page 1961

Procedure

Step 1 Click **Add Application Group**.

Step 2 In the **Application Group** section, type the name in the **Name** field and click **Next**.

Step 3 In the **SAML Service Provider (SP) Metadata** section, the data is dynamically generated. Copy the values of the **Entity ID** and **Assertion Consumer Service (ACS) URL** fields or click **Download SP Metadata** to download this data in XML format for adding it to the IdP. Click **Next**.

Step 4 In the **SAML Identity Provider (IdP) Metadata** section, add the metadata using any one of the methods:

- **XML File Upload**—Choose a file or drag and drop the XML file.

  The details of the **Entity ID**, **Single Sign-On URL**, and **IdP Certificate** are displayed.

- **Manual Configuration**—Perform these steps:
  - **Entity ID**—Enter the URL that is defined in the SAML IdP to identify a service provider uniquely.
  - **Single Sign-On URL**—Enter the URL for signing into the SAML identity provider server.
  - **IdP Certificate**—Choose the certificate of the IdP enrolled in threat defense to verify the messages signed by the IdP.

Click the **Add** (+) icon to configure a new certificate enrollment object. For more information, see Add certificate enrollment, on page 1964.

- **Configure Later**—In the event you do not have the IdP metadata, you can configure it later.

Click **Next**.

Step 5 In the **Re-authentication Interval** section, enter the value in the **Timeout Interval** field and click **Next**.

The re-authentication interval allows you to provide a value that determines when a user must authenticate again.

Step 6 In the **Security Zones and Security Controls** section, the security zones and threat settings are inherited from the parent policy. You can override these settings. Click **Next**.

Step 7 Review the configuration summary. Click **Edit** to modify the details in any of the sections. Click **Finish**.

Step 8 Click **Save**.

The Application Group is created and is displayed on the Zero Trust Application page.

What to do next

1. Create an application, on page 1965.
2. Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Add certificate enrollment

This task allows you to add certificate enrollment by configuring the certificate name and installing the IdP certificate in PEM format.

Certificate enrollment is required to establish secure authentication. You need to provide certificate information in the correct format to ensure proper system configuration.

Procedure

Step 1 Enter the **Name**.

Step 2 Paste the certificate information in the **IdP Certificate** field in PEM format.

Note
If the certificate is dependent on a root or intermediate certificate, you must install the dependant certificates. See Certificates, on page 3079.

Step 3 Click **Save**.

Create an application

This task creates a Grouped or Ungrouped Application to establish secure access controls for private applications within your Zero Trust Access policy framework.

Before you begin

1. Create a zero trust application policy, on page 1961.
2. Create an application group, on page 1963 (required only for Grouped Applications).

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**

Step 2 Choose the policy.

Step 3 Click **Add Application**.

Step 4 In the **Application Settings** section, complete the following fields.

- **Application Name**—Enter the application name.
- **External URL**—Enter the URL that is used by the user to access the application.
- **Application URL**—By default, the external URL is used as the Application URL. Uncheck the **Use External URL as Application URL** check box to specify a different URL.

  If Threat Defense uses an internal DNS, the Application URL must match an entry within that DNS to ensure it resolves to the application. In addition, if you enter a fully qualified domain name (FQDN) that could have either an IPv4 or IPv6 address, the FQDN must resolve to both IPv4 and IPv6. Threat Defense does not translate between these types of addresses.

- **Application Certificate**—Choose the certificate for the private application. Click the **Add** (+) icon to configure an internal certificate object. For more information, see Add internal certificate objects, on page 2997.
- **Source Network Address Translation**—Source Network Address Translation (source NAT) translates the source IP address in a packet to an address from a specified range.

  Enable **Source Network Address Translation** to specify an IP address range. The source IP address of an incoming request is translated to an IP address from the specified range. Source NAT works for both IPv4 and IPv6 addresses.

  - **IPv4**—Choose the source network for NAT from the list. Click **Add** (+) to create a network object. For more information, see Networks, on page 2982.
  - **IPv6**—Choose the source network for NAT from the list. Click **Add** (+) to create a network object. For more information, see Networks, on page 2982.

  Source NAT translates the source IP address of an incoming request to a routable IP address specified in the network object or the object group.

  Note

Only object or object groups of type Host or Range are supported.

- **Application Group**—Choose the Application Group from the drop-down list. See Create an application group, on page 1963.

  Note
  This field is not applicable for an ungrouped application.

Step 5 Click **Next**.

Step 6 Depending on the type of application:

- For a Grouped Application, the **SAML Service Provider (SP) Metadata**, **SAML Identity Provider (IdP) Metadata**, and **Re-authentication Interval** are inherited from the Application Group and do not need to be configured by the user.
- For an Ungrouped Application, perform these steps:

  a. In the **SAML Service Provider (SP) Metadata** section, the data is dynamically generated. Copy the **Entity ID**, **Assertion Consumer Service (ACS) URL** of the IdP, or click **Download SP Metadata** to download this data in XML format for adding it to the IdP. Click **Next**.

  b. In the **SAML Identity Provider (IdP) Metadata** section, add the metadata using any one of the methods:

  - **XML File Upload**—Choose a file or drag and drop the XML file.

    The details of the **Entity ID**, **Single Sign-On URL**, and **IdP Certificate** are displayed.

  - **Manual Configuration**—Perform these steps:

    - **Entity ID**—Enter the URL that is defined in the SAML IdP to identify a service provider uniquely.
    - **Single Sign-On URL**—Enter the URL for signing into the SAML identity provider server.
    - **IdP Certificate**—Choose the certificate of the IdP enrolled in threat defense to verify the messages signed by the IdP.

      Click the **Add** (+) icon to configure a new certificate enrollment object. For more information, see Add certificate enrollment, on page 1964.

  - **Configure Later**—If you do not have the IdP metadata, you can configure it later.

  Click **Next**.

  c. In the **Re-authentication Interval** section, enter the value in the **Timeout Interval** field and click **Next**. The reauthentication interval allows you to provide a value that determines when a user must authenticate again.

Step 7 In the **Security Zones and Security Controls** section, the security zones and threat settings are inherited from the parent policy or application group. You can override these settings. Click **Next**.

Step 8 Review the configuration summary. Click **Edit** to modify the details in any of the sections. Click **Finish**.

Step 9 Click **Save**.

The application is listed on the Zero Trust Application page and enabled by default.

What to do next

1. Set targeted devices for zero trust access policy, on page 1968.
2. Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Create an ungrouped application

This task creates a standalone application for zero trust application access.

Before you begin

Create a zero trust application policy, on page 1961

Procedure

Step 1 Choose **Policies** > **Access Control** > **Zero Trust Application**

Step 2 Select the policy and click **Add Application**.

Step 3 In the **Application Settings** section, complete the following fields.

- **Application Name**—Enter the application name.
- **Application Group**—This field is not applicable for a standalone application.
- **External URL**—Enter the URL that is used by the user to access the application.
- **Internal URL (FQDN or Network IP)**—Enter the URL where the application is hosted on the internal server.
- **Application Certificate**—Enter the certificate for the private application.

Click **Next**.

Step 4 In the **SAML Service Provider (SP) Metadata** section, the data is dynamically generated. Copy the Entity ID or ACS URL of the IdP or click **Download SP Metadata** to download this data in XML format for adding it to the IdP. Click **Next**.

Step 5 In the **SAML Identity Provider (IdP) Metadata** section, add the metadata using either one of the methods:

- **XML File Upload**—Select a file or drag and drop the XML file.
- **Manual Configuration**—Complete the following fields:
  - **Entitv ID**—The URL that is defined in the SAML IdP to identify a service provider uniquely.
  - **Single Sign-On URL**—The URL for signing into the SAML identity provider server.
  - **Certificate**—Certificate of the IdP enrolled into the threat defense to verify the messages signed by the IdP.

- **Configure Later**—In the event you do not have the IdP metadata, you can configure it later.

Click **Next**.

Step 6 In the **Re-authentication Interval** section, enter the value in the **Interval** field and click **Next**.

Step 7 In the **Interface Access and Security Controls** section, the interface objects and threat settings are inherited from the parent policy. You can override these settings.

Step 8 Review the settings and click **Finish**.

Step 9 Click **Save**.

The Application is displayed on the policy page under **Ungrouped Applications**.

What to do next

Set targeted devices for zero trust access policy, on page 1968

Set targeted devices for zero trust access policy

Set the specific devices where a Zero Trust Application policy will be deployed to control application access and enforce security policies.

Each Zero Trust Application policy can target multiple devices; each device can have one deployed policy at a time.

Before you begin

1. Create a zero trust application policy, on page 1961.
2. Create an application group, on page 1963.
3. Create an application, on page 1965.

Follow these steps to set targeted devices for Zero Trust Access Policy:

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**

Step 2 Choose the policy.

Step 3 Click **Targeted Devices**.

Step 4 Choose the devices where you want to deploy the policy using any one of the methods:

- Choose a device in the **Available Devices** list and click >> or the **Add** (⊕) icon.
- To remove a device from the **Selected Devices** list, choose a device and click << or the **Delete** (🗑) icon.

Step 5 Click **Apply** to save policy assignments.

Step 6 Click **Save** to save the policy.

What to do next

Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Edit a zero trust application policy

Edit the settings of a Zero Trust Application Policy to modify policy configurations, application groups, or individual applications according to your security requirements.

Follow these steps to edit the settings of a Zero Trust Application Policy.

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**

Step 2 Click **Edit** (✎) next to the Zero Trust Application Policy you want to edit.

Step 3 Edit your Zero Trust Application Policy.

You can change these settings or perform these actions:

- Name and Description—Click **Edit** (✎) next to the policy name, make your changes, and click **Apply**.
- To modify the policy settings:
  - Click **Settings**
  - Modify the settings as required.

    Important
    Editing the domain name for the SAML ACS URL interrupts application access.

  - Click **Save**.
- To modify the Application Group settings:
  - Click **Applications**.
  - Click **Edit** (✎) next to the Application Group you want to edit.
  - In each section, click **Edit** to modify the settings, as required

    Important
    Editing the Application Group name interrupts application access.

  - Click **Apply** after you modify the settings in a section.
  - Click **Finish**.
  - Click **Save**.

- To modify the Application settings:
  - Click **Applications**.
  - Click **Edit** (✎) next to the Application you want to edit.
  - In each section, click **Edit** to modify the settings, as required.

    Important
    Editing the Application name interrupts application access.

  - Click **Apply** after you modify the settings in a section.
  - Click **Finish**.
  - Click **Save**.
- To enable, disable, or delete multiple Applications, choose the Applications, click the required bulk action, and click **Save**.

  Note
  These actions are also available in the right-click menu.

  - To enable all Applications, click **Bulk Actions** > **Enable**.
  - To disable all Applications, click **Bulk Actions** > **Disable**.
  - To delete all Applications, click **Bulk Actions** > **Delete**.
- Click **Return to Zero Trust Application** to return to the policy page.

The Zero Trust Application Policy settings are updated with your changes.

What to do next

Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Manage zero trust application policies

This task allows you to manage zero trust application policies to control access and security for applications in your network.

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**

Step 2 Choose the appropriate action for the zero trust application policies.

- Create—Click **New Policy**. See Create a zero trust application policy, on page 1961

- Edit—Click **Edit** ( ). See Edit a zero trust application policy, on page 1969
- Report—Click **Report** ( ).
- Delete—Click **Delete** ( ).

Step 3 Click **Save**.

What to do next

Ensure that there are no warnings before you deploy the configuration to threat defense. To deploy configuration changes, see Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Monitor zero trust sessions

This task enables administrators to effectively monitor Zero Trust Application Policy sessions through multiple methods including connection events, dashboard visualization, CLI commands, and diagnostic troubleshooting tools.

After a Zero Trust Application Policy is deployed, additional monitoring capabilities become available to track session activity, user behavior, and system performance. To monitor the zero trust sessions, do these steps.

Procedure

Step 1 Add Zero Trust fields to the connection events table view.

a. Choose **Analysis** > **Connections** > **Events**.

b. Click the **Table View of Connection Events** tab.

c. In the table view of events, multiple fields are hidden by default. To change the fields that are displayed, click the **x** icon in any column name to display a field selector.

d. Select these fields:

- **Authentication Source**
- **Zero Trust Application**
- **Zero Trust Application Group**
- **Zero Trust Application Policy**
- **Zero Trust Application Host**
- **Zero Trust Origin User**
- **Zero Trust Proxy**

- **Zero Trust Rule**
- **Zero Trust Status**
- **Zero Trust Tunnel ID**

e. Click **Apply**.

See Connection and Security-Related Connection Events in the Secure Firewall Management Center Administration Guide for more information on the connection events.

Step 2 Access the Zero Trust dashboard to monitor real-time session data.

The Zero Trust dashboard provides a summary of the top zero trust applications and zero trust users that are managed by the management center.

Choose **Overview** > **Dashboards heading** > **Dashboard**, and click the **Zero Trust** tab to access the dashboard.

The dashboard has the following widgets:

- Top Zero Trust Applications
- Top Zero Trust Users

Step 3 Use CLI commands to monitor and troubleshoot Zero Trust configurations.

Log in to the device CLI and use the following commands:

CLI Command	Description
show running-config zero-trust	View the running configuration for a zero trust configuration.
show running-config zero-trust-hybrid	View the running configuration for a universal zero trust configuration.
show zero-trust	Display the run-time zero trust statistics and session information.
show cluster zero-trust	Display the summary of zero trust statistics across nodes in a cluster.
clear zero-trust	Clear zero trust sessions and statistics.
show counters protocol zero_trust	View the counters that are hit for zero trust flow.

Step 4 Run diagnostic tools to troubleshoot Zero Trust configuration issues.

The diagnostics tool facilitates the troubleshooting process by detecting possible issues with zero trust configurations. The diagnostics can be classified into two types:

- **Application-specific diagnostics** are used to detect issues such as:
  - DNS-related issues
  - Misconfigurations such as socket not open, and issues with classification and NAT rules.
  - Issues with deployment of zero trust policy or SSL rules

    - Issues with source NAT issues and exhaustion of PAT pool

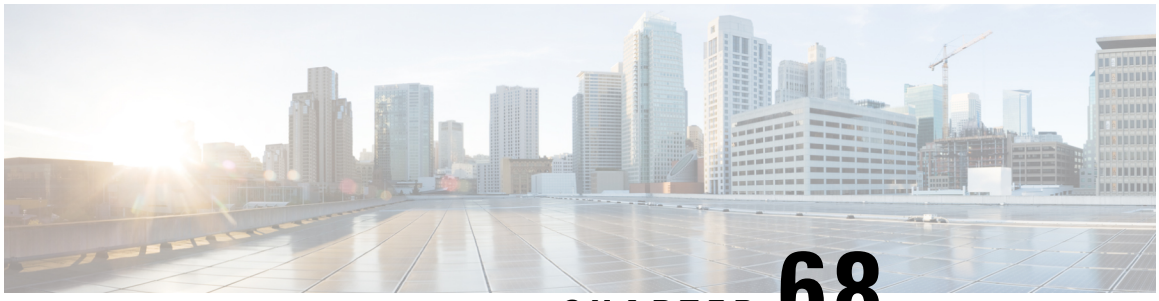
- **General diagnostics** are used to detect issues such as:
    - Strong cipher license not enabled
    - Invalid application certificate
    - SAML-related issues
    - Home agent and cluster bulk sync issues

a. Click Diagnostics ( ) next to the zero trust application that you want to troubleshoot. The **Diagnostics** dialog box appears.

b. Choose the device from the **Select Device** drop-down list and click **Run**. A report is generated in the **Reports** tab after the diagnostic process is complete.

c. To view, copy, or download the logs, click the **Logs** tab.

You have successfully configured comprehensive monitoring for Zero Trust sessions. The connection events table now displays Zero Trust-specific fields, the dashboard provides real-time visualization of top applications and users, CLI commands are available for detailed system monitoring, and diagnostic tools can identify and troubleshoot configuration issues.



CHAPTER 68

Universal Zero Trust Network Access

Universal Zero Trust Network Access (ZTNA) secures internal resources by validating user identity and device posture instead of granting full network access. The solution integrates Security Cloud Control, Secure Access, and Firewall Threat Defense devices to enforce granular, context-aware security policies. Administrators configure these components within the Security Cloud Control platform to protect applications and prevent lateral movement across hybrid environments.

- Universal Zero Trust Network Access, on page 1975
- Threat Defense with Universal ZTNA, on page 1976
- Prerequisites for Universal Zero Trust Network Access, on page 1977
- Restrictions for Universal Zero Trust Network Access, on page 1978
- Enable Cloud-Delivered Firewall Management Center in security cloud control, on page 1979
- Configure Security Devices, on page 1979
- Configure network connections, on page 1985
- Related documentation, on page 1993

Universal Zero Trust Network Access

Universal Zero Trust Network Access (universal ZTNA) is a client-based ZTNA solution that enables users to securely access internal resources and applications regardless of their location, whether remote or on-premises. It enables administrators to specifically allow access to internal network resources according to user identity including user trust and posture, without granting access to the entire network as with Remote Access VPN.

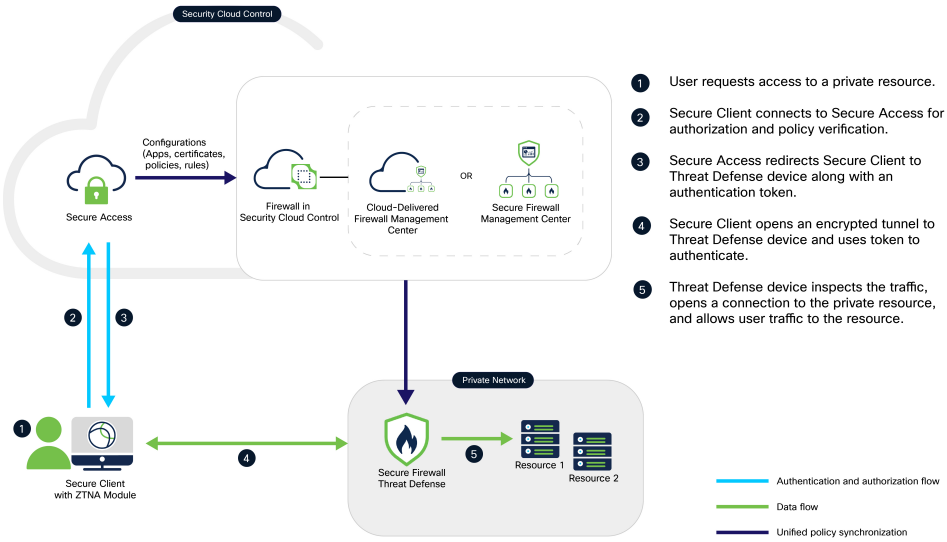
Because universal ZTNA does not assume that access granted to one application implicitly authorizes access to other applications, the network attack surface is reduced.

Universal ZTNA ensures least-privileged, per-application, per-user access with strong authentication, posture validation, and comprehensive traffic inspection. It secures applications effectively across hybrid environments.

Components of Universal ZTNA

A new configuration of universal ZTNA consists of Security Cloud Control Firewall Management (formerly called Cisco Defense Orchestrator), and Secure Access, both provisioned on the Security Cloud Control platform. Security Cloud Control Firewall Management manages the Firewall Threat Defense devices through the Secure Firewall Management Center.

Figure 381: Components of universal ZTNA



- **Security Cloud Control Firewall Management**: Manages the configuration and deployment of universal ZTNA policies to the Firewall Threat Defense devices. The Threat Defense devices protect on-premises resources by enforcing universal ZTNA policies. Threat Defense inspects traffic and enforces intrusion prevention system (IPS), file, and malware policies on the traffic.
- **Secure Access**: Secure Access defines the access policies, posture, and security profiles for the user. It enforces the policies for user traffic through the cloud.
- **Security Cloud Control platform**: Security Cloud Control provides a unified secure management plane for both Secure Access and Firewall, simplifying the administration of universal ZTNA policies across them.
- **Secure Client**: The Secure Client is installed on the end user's device. It acts as the enforcement point that intercepts connection requests to protected internal resources, enabling secure, identity-based access.

The following sections describe how to enable universal ZTNA on a Threat Defense device. For the complete configuration of universal ZTNA, refer to the *Universal Zero Trust Network Access Configuration Guide*.

Threat Defense with Universal ZTNA

A universal ZTNA-enabled Firewall Threat Defense device is a critical enforcement point that protects private resources by integrating zero trust principles with firewall capabilities. Here's how it operates to secure private resources:

- Enforces access and validates policy: A Threat Defense device intercepts all access requests to private resources. It enforces granular, context-aware access policies that verify user identity, device posture, and contextual factors before granting access. Only requests that meet the least-privilege criteria are allowed, blocking the unauthorized attempts to access resources.
- Establishes a secure tunnel to the resource: After validating the access policy, the Threat Defense device establishes a secure, encrypted tunnel between the user's device and the private resource. This tunnel ensures that private resources are isolated from direct network exposure.

- Inspects traffic: Threat Defense device continuously inspects the traffic and blocks threats before they reach the resources. The device applies the intrusion prevention system, file, and malware detection capabilities to detect and block the threats.
- Microsegments and prevents lateral movement: The device enforces microsegmentation by restricting traffic flows to only the authorized resources. This containment prevents lateral movement within the network, limiting the impact of any potential breach.

Prerequisites for Universal Zero Trust Network Access

This topic discusses requirements and guidelines for Universal Zero Trust Network Access (universal ZTNA).

Licensing requirements

- Secure Firewall Management Center requires a smart license account with export-controlled features. It does not function in evaluation mode for universal ZTNA.

  Secure Firewall Threat Defense devices require an IPS license if Intrusion policy is configured. If Malware policies are configured, the devices require IPS and Malware Defense licenses. For more information, refer to the "Licenses" section in the *Cisco Secure Firewall Management Center Administration Guide*.
- Secure Access requires a subscription of Cisco Secure Private Access Essentials or Advantage.

Device requirements

- All Cloud-Delivered Firewall Management Center and Secure Firewall Threat Defense devices must be running Version 7.7.10 or later.
- All Secure Firewall Threat Defense devices must be configured for routed mode; transparent mode is not supported.
- In Security Cloud Control, when you are configuring universal zero trust access for a device, ensure that the Enrollment Type for the device identity certificate is an object that is created using the PKCS12 file format. No other certificate type is supported. If necessary, you can also create a new certificate object from Security Cloud Control, which supports the PKCS12 format. See Configure Security Devices.
- Configure the Domain Name System (DNS) to resolve Fully Qualified Domain Name (FQDN) of private resources. Use the Platform Settings menu on the Secure Firewall to configure the DNS. See Interface and Device Settings.
- High Availability (HA) devices *are* supported; they are displayed as one entity.
- Secure Client (with ZTNA module enabled) Version 5.1.10 and later is supported.

  The client must be running in a platform that supports Trusted Platform Module (TPM), such as Windows 11.

Guidelines on certificate types

- **User Identity Certificate**: Secure Client, which is zero trust access enabled, presents the user identity certificate during the Mutual Transport Layer Security (mTLS) session with Secure Access and Firewall Threat Defense to request access to private resources.

- **Firewall Threat Defense Device Certificate**: Threat Defense devices that are universal ZTNA-enabled use device certificates to establish secure mTLS connections with the Secure Client and Secure Access. Ensure that the device identity certificate is of type PKCS12.

  If you have already enrolled a manual certificate for the device, first export it to the PKCS12 format using the **Devices** > **Certificates** > **Export Certificate** menu on Firewall Management Center. Use the exported PKCS12 file to create a new PKCS12 certificate enrollment object.

- **Decryption Certificate**: (Optional) To decrypt the traffic that is sent to private resources, enable **Decryption** for the resources in Secure Access and provide the server certificate and key. We recommend that you use a certificate that is signed by a publicly recognized certificate authority (CA).

Supported devices

Both on-premises Firewall Management Center and Cloud-Delivered Firewall Management Center can be configured to manage the devices.

Only devices that have 16 cores or more are supported. Such models of Secure Firewall Threat Defense

- 1150
- 3105, 3110, 3120, 3130, 3140
- 4115, 4125, 4145, 4112
- 4215, 4225, 4245
- FTDv

Note You can configure universal ZTNA on a 16 core FTDv. No other deployment configuration of FTDv supports universal ZTNA.

Restrictions for Universal Zero Trust Network Access

Be aware that Universal ZTNA has several operational limitations that may affect deployment and functionality.

- Universal ZTNA does not support IPv6.
- Universal ZTNA-enabled devices do not enforce policies for traffic over a site-to-site tunnel.
- Universal ZTNA does not support clustered devices.
- Universal ZTNA does not support devices in multi-instance mode.
- Universal ZTNA sessions do not support jumbo frames.
- Universal ZTNA supports only global VRF.
- Universal ZTNA does not support protocols such as FTP or TFP, where the data or secondary connection originates from a server.

For example, an active FTP connection uses a persistent control connection for commands and creates temporary data connections for file transfers. Universal ZTNA does not support such data connections that originate from the server.

Enable Cloud-Delivered Firewall Management Center in security cloud control

Enable the Cloud-Delivered Firewall Management Center to configure universal ZTNA and manage associated Secure Firewall Threat Defense devices.

To configure universal ZTNA, first enable the Cloud-Delivered Firewall Management Center in Cisco Security Cloud Control. After you enable it, configure the associated Secure Firewall Threat Defense devices. The devices can then use the networks, private resources, and policies needed to configure and manage universal ZTNA.

Procedure

Step 1 From the Security Cloud Control menu, click **Administration** > **Integrations** > **Firewall Management Center**.

Step 2 Click **Enable Cloud-Delivered FMC**.

Security Cloud Control starts provisioning the instance. This process takes a few minutes to complete. You can track the progress in the **Status** column of **Cloud-Delivered FMC**.

After the provisioning is complete, the status changes to **Active**.

Onboard the Firewall Threat Defense devices to the Cloud-Delivered Firewall Management Center. To learn how to onboard a Threat Defense device, see Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center

The Cloud-Delivered Firewall Management Center is enabled and active in Security Cloud Control, ready for Secure Firewall Threat Defense device configuration.

Configure Security Devices

All Firewall Threat Defense devices associated with the Cloud-Delivered Firewall Management Center that you onboarded to Security Cloud Control are *security devices* to which you can:

- Associate private resources, which are internal applications you want to protect with identity-based access control, IPS, malware, and other protections.
- Deploy Secure Access access rules. Security devices are responsible for enforcing access rules for on-premises users, remote users, or both.

Perform these steps to enable **universal zero trust network access settings** on the Threat Defense devices. These steps include configuring the device FQDN, inside interface, outside interface, and PKCS12 certificate to enable universal ZTNA on the devices.

Before you begin

You must know the name of each device's internal and external network interfaces:

- The internal interface (also referred to as the *DMZ* interface) is used to apply access rules to on-premises users.
- The external interface is used to apply access rules to remote users.

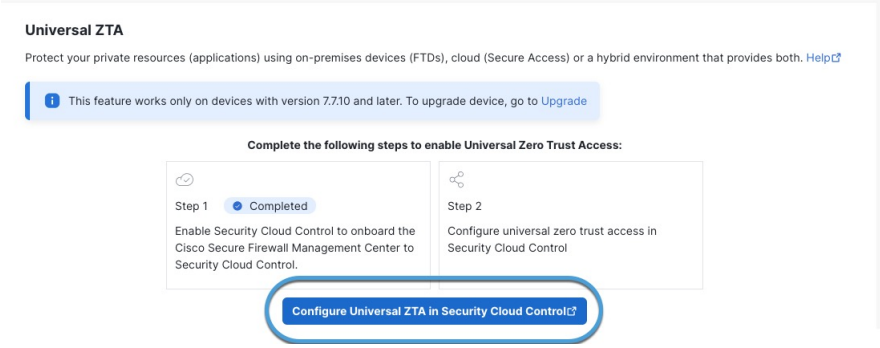
You can choose internal, external, or both types of interfaces for each security device.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, click **Policies** > **Zero Trust Application**.

Step 2 Click **Configure Universal ZTA in Security Cloud Control**.

This figure shows an example.

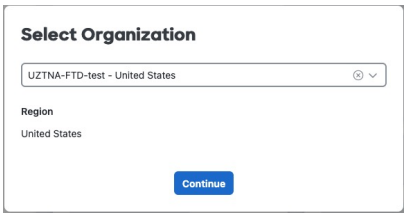


Step 3 When prompted, log in to Security Cloud Control.

Step 4 When prompted, select your organization from the drop-down list and click **Continue**.

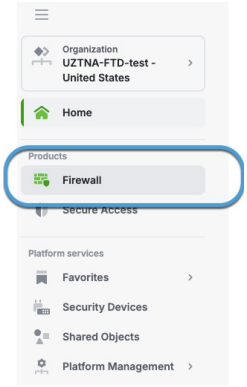
Select an organization that has both Secure Access and Secure Firewall micro applications configured.

This figure shows an example.



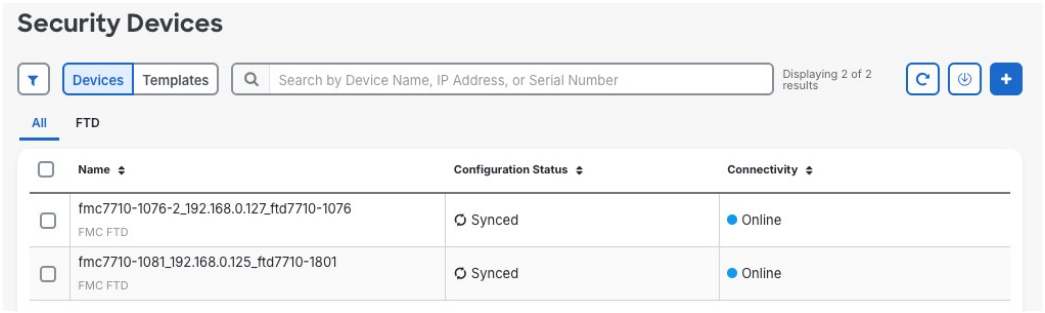
Step 5 In Security Cloud Control, in the Products section, click **Firewall**.

This figure shows an example.



Step 6 In the Manage section, click **Security Devices**.

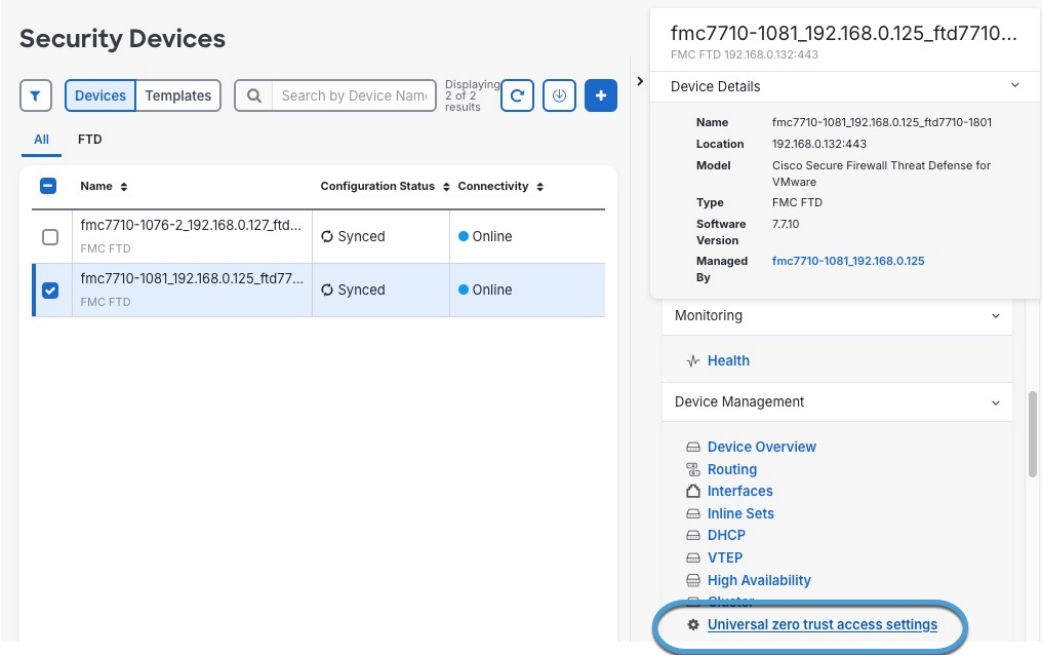
The **Security Devices** page displays the available security devices.



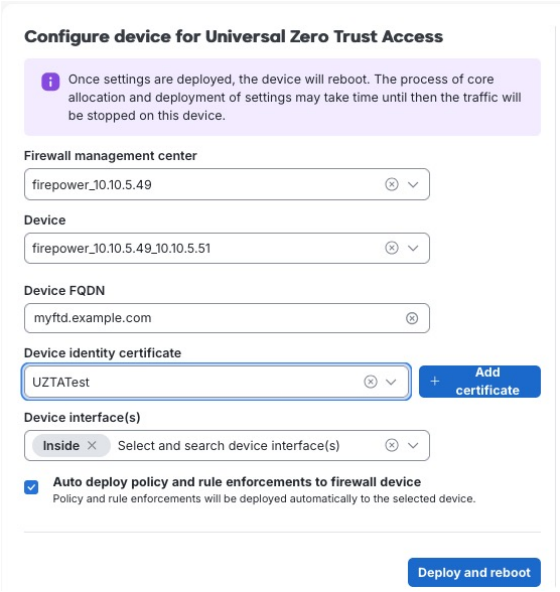
Step 7 Select the check box next to a device to add to the universal zero trust network access configuration.

Step 8 In the right pane, click **Device Management** > **Universal zero trust access settings**.

This figure shows an example.



Step 9 Enter or edit the following information on the **Configure device for Universal Zero Trust Access** page.



This table describes the configurations to enable universal ZTNA on the device.

Item	Description
Firewall management center	From the drop-down list, click the name of a Cloud-Delivered Firewall Management Center to use for policy deployment, monitoring, and other tasks.
Device	From the drop-down list, click the name of a device to use for rule deployment and enforcement.

Item	Description
Device FQDN	Enter the security device's fully qualified domain name (FQDN). The FQDN is also referred to as the TLS/SSL certificate's Common Name. Secure Access redirects the clients to the resource that is represented by this FQDN. Ensure that the device FQDN *exactly matches* the Common Name (CN) of the device identity certificate and must also match a Subject Alternative Name (SAN) in the certificate.
Device identity certificate	The **Device identity certificate** must have a Common Name that: • *Exactly matches* the value you enter in the **Device FQDN** field. • Matches a Subject Alternative Name (SAN) in the certificate. From the drop-down list, click the name of an existing identity certificate from the list. Click **Add certificate** and add an identity certificate in **.p12** format (also referred to as PKCS#12; see this article on ssl.com). **Note** Universal ZTNA supports only the PKCS#12 format of certificate enrollment. In the provided fields, enter a **Name** to identify the certificate. Then copy/paste, drag/drop, or upload the certificate and private key. If the certificate is encrypted, enter its password in the provided field. You can optionally use a wildcard certificate as discussed in What is a Wildcard Certificate? on ssl.com.
Device Interface(s)	From the drop-down list, select the check box next to any of the following types of interfaces. • Internal network interface (or DMZ): deploys access rules for on-premises users only. • External network interface: deploys access rules for remote users only. • Both types of interfaces: deploys access rules for either on-premises or remote users.
Auto deploy policy and rule enforcements to firewall device	Select the check box to automatically deploy access rules to the device after they are updated on Secure Access. On the device, the Auto deploy feature selectively deploys only the Universal ZTNA access policy. It does not impact other changes or configurations on the Firewall Management Center. **Note** If there are other interdependent policies on the device (which are interlinked with the Universal ZTNA access policy), the Firewall configuration status displays an error message. The deployment then stops. In such cases, you should manually deploy the Universal ZTNA access policy from the Firewall Management Center.

Step 10 Click **Deploy and Reboot**.

The device reboots to reallocate the system resources for universal ZTNA components.

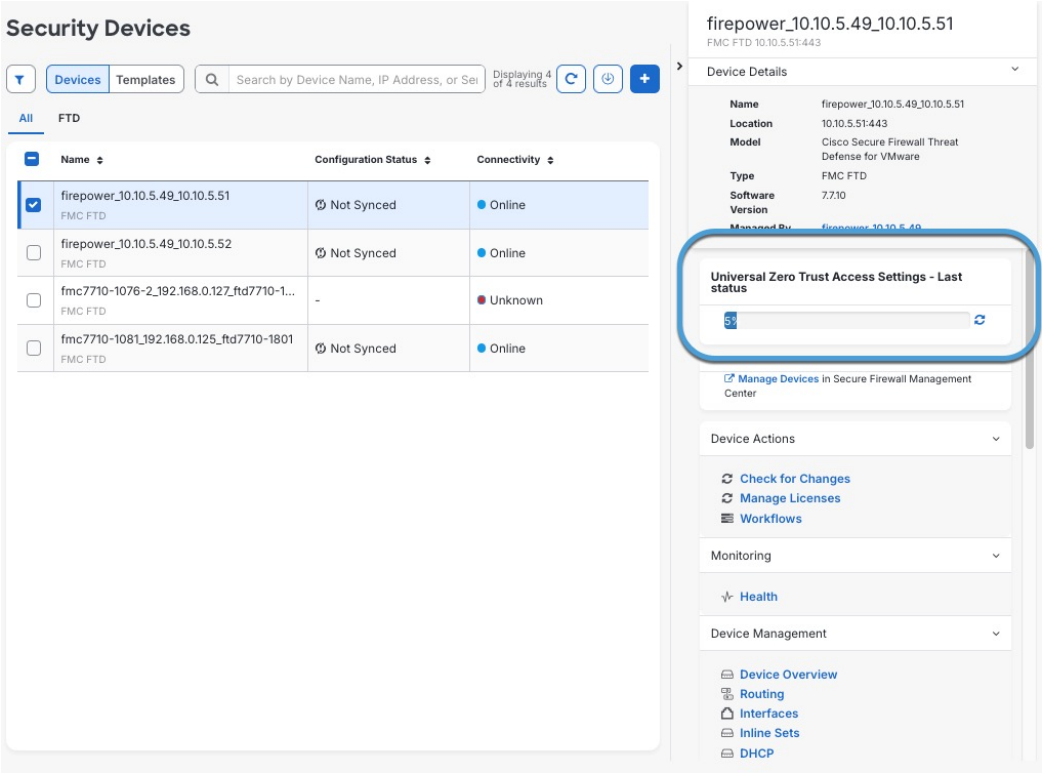
Note
The device takes several minutes to reboot, during which time all traffic handled by the device is disrupted.

If you deploy a High Availability (HA) pair of devices, both devices reboot simultaneously.

This behavior is applicable to only 10.0 version.

Step 11 On the **Security Devices** page, select the check box next to the device to which you just deployed the Universal ZTNA configuration.

The right pane displays the deployment status, as shown in the figure.



For additional information, click **Device Actions** > **Workflows** in the right pane.

After the deployment completes, you can view the completion status in the **Universal Zero trust Access Settings - Last status** tab for the device.

Universal ZTNA-enabled Firewall Threat Defense device is connected to Secure Access.

What to do next:

Check the availability of the Threat Defense device under Secure Access by clicking **Security Cloud Control** > **Secure Access** > **Connect** > **Network Connections** > **FTD**.

Configure network connections

Configure private resources and the devices that protect these resources to complete the Universal ZTNA configuration.

The final step in configuring Universal ZTNA is to configure private resources and the devices that are responsible for protecting the resources.

Before you begin

Complete the tasks discussed in Configure Security Devices, on page 1979.

Procedure

Step 1 In Security Cloud Control, click **Products** > **Secure Access**.

The Secure Access product menu appears in the left navigation bar.

Step 2 Click **Resources** > **Destinations** > **Private Resources**.

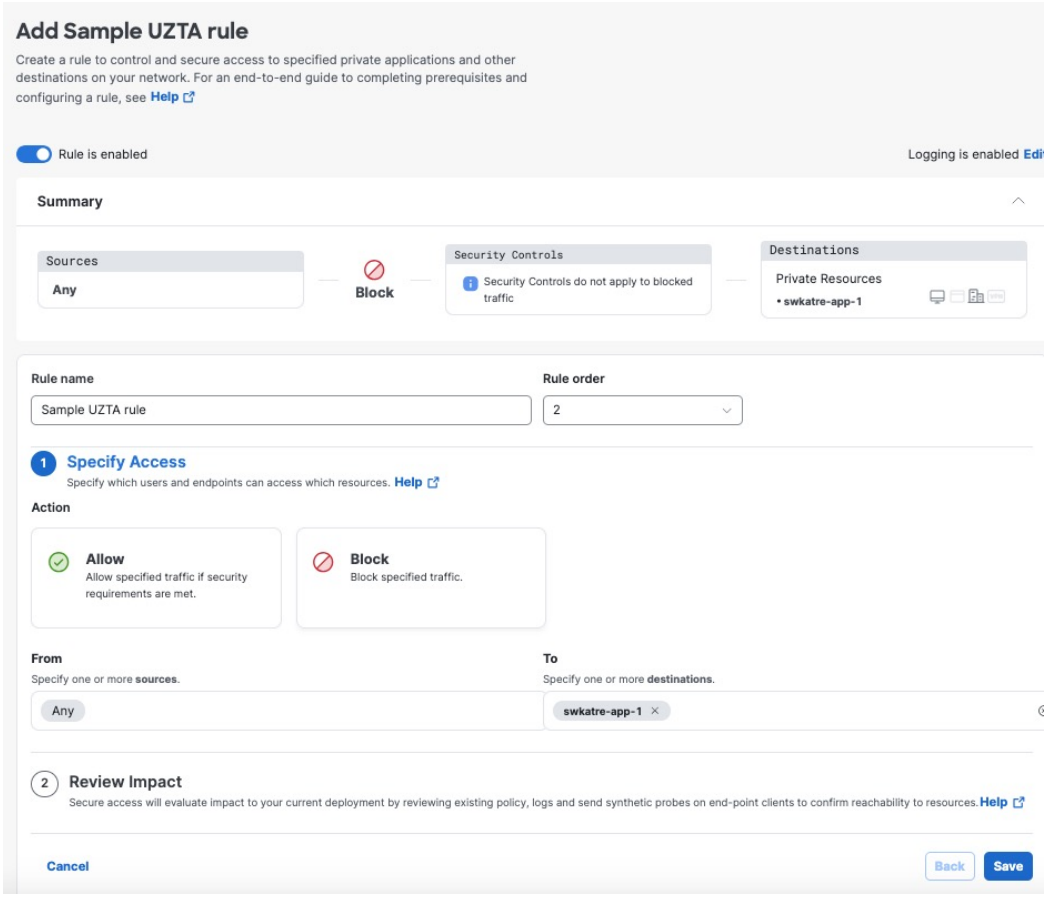
Add private resources. For guidance, refer to Manage Private Resources.

Note
When you add a private resource, select the zero-trust connections method of endpoint connections. This selection enables client-based zero trust access for the private resource.

Step 3 Click **Secure** > **Access Policy**.

Add or edit access rules. For more information, refer to Manage the Access Policy.

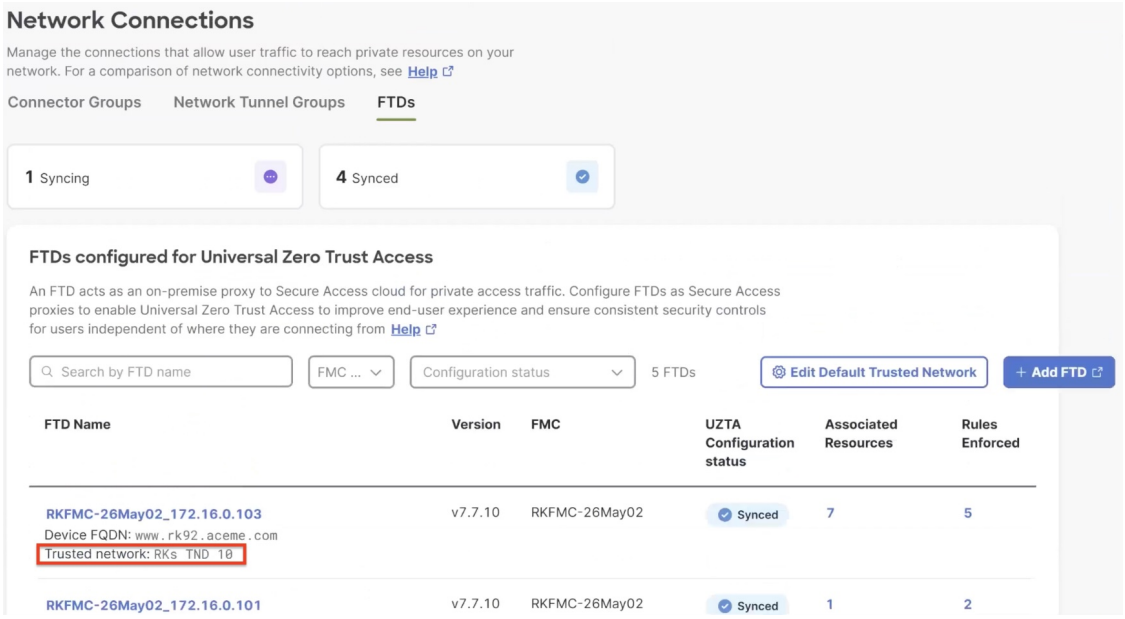
This sample access rule blocks access to a destination named **swatw-app-1**:



Step 4 Click **Connect** > **Network Connections**.

Step 5 Click the **FTDs** tab.

The page displays the available Secure Firewall Threat Defense devices configured for universal zero trust network access.



Before proceeding to the next step, ensure that the device is associated with a trusted network. This association enforces policies on traffic originating from the trusted network.

After onboarding a Threat Defense device, it is automatically associated with a default trusted network if one exists. If not, you must create a trusted network and associate it with the Threat Defense device.

Step 6 Click the name of a Threat Defense device to configure.

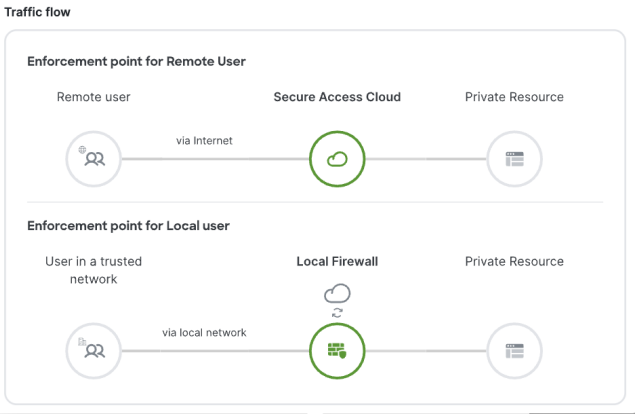
Step 7 In the right pane, click **Associate Resources**.

Note

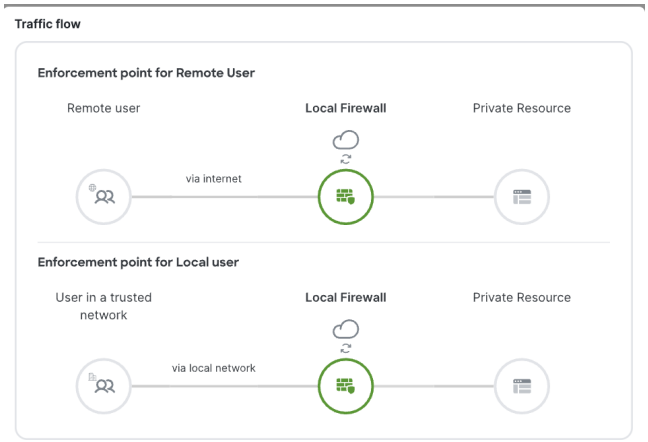
- Only those resources that are enabled for zero trust access can be associated with a Threat Defense device.
- A Threat Defense device must have connectivity to the associated private resources.
- Resources associated with a Threat Defense device are shared with other devices that have the same FQDN.

Step 8 In the **Associate Private Resources** dialog box, make the following selections to specify the access policy enforcement and traffic flow for a user:

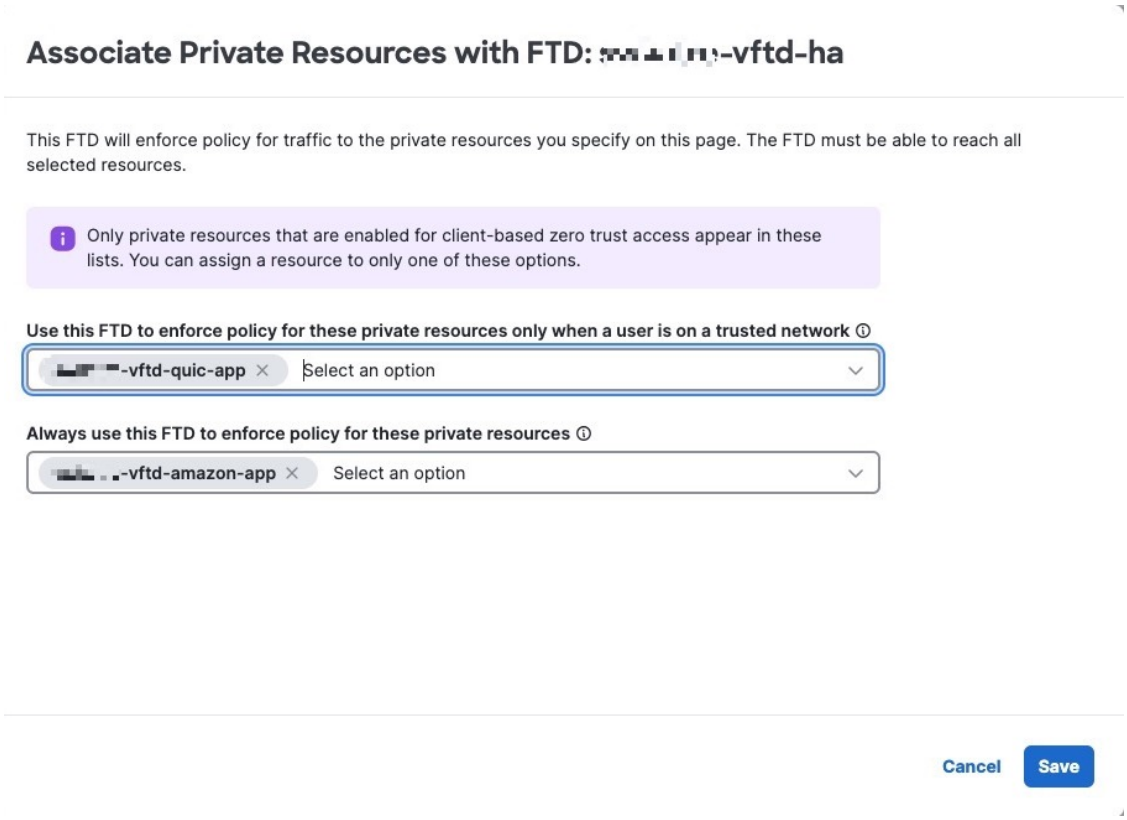
- **Use Threat Defense device to enforce policy only for on-premises users**: From the **Use this FTD to enforce policy** drop-down list, select the private resources that a user should access only from an on-premises location.



- **Use Threat Defense device to enforce policy for both on-premises and remote users**: From the **Always use this FTD to enforce policy** drop-down list, select the private resources for which the selected Threat Defense device always enforces policy, regardless of whether the user is on-premises or remote.



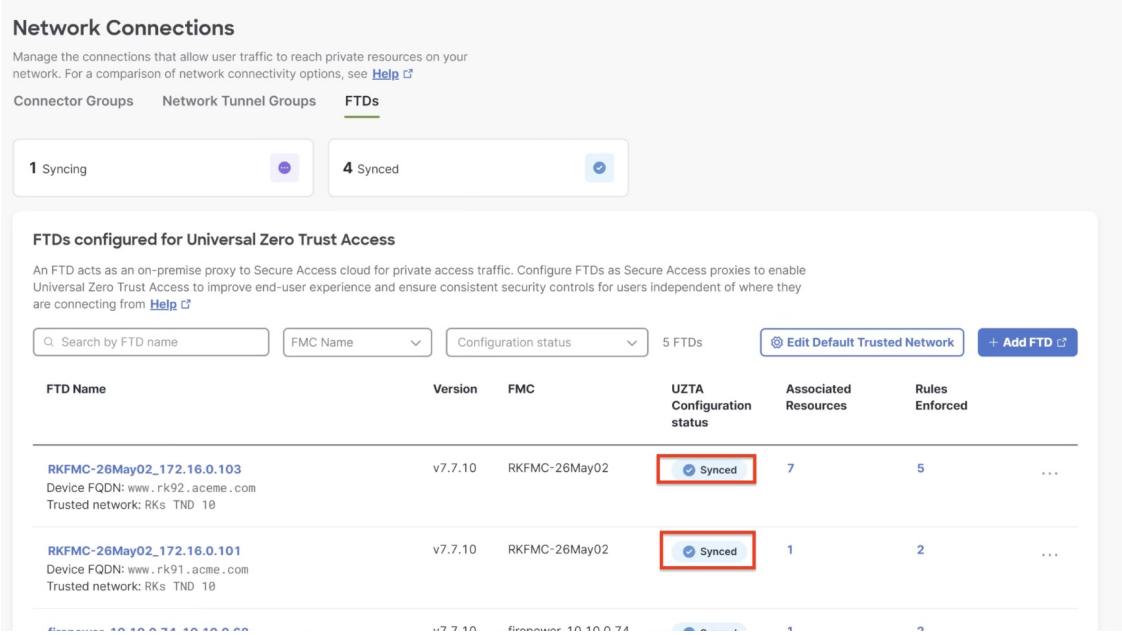
The following figure shows an example of using a Threat Defense device to enforce access rules for the **vftd-quic-app** for on-premises users and **vftd-amazon-app** for all users, whether on-premises or remote.



Step 9 Click **Save**.

The configurations are applied to the device, and the **UZTA Configuration status** column for the device displays **Synced**.

The following figure shows an example.



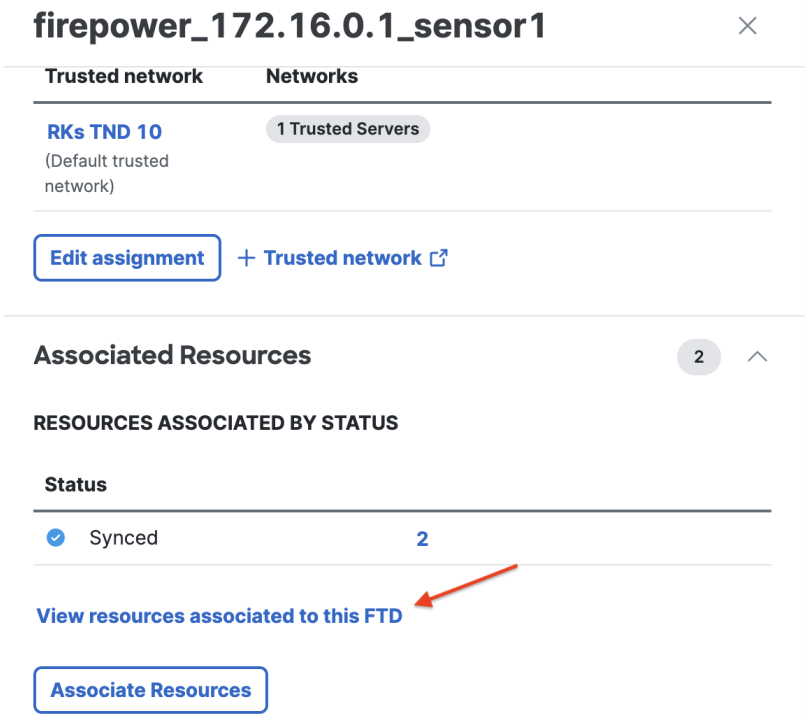
Configuration status can also be:

- Syncing—updates to the Threat Defense device are ongoing.
- Out of sync—modifications to Secure Access configurations are pending update to the Threat Defense device.
- Failed to sync—configurations were not updated on the Threat Defense device.

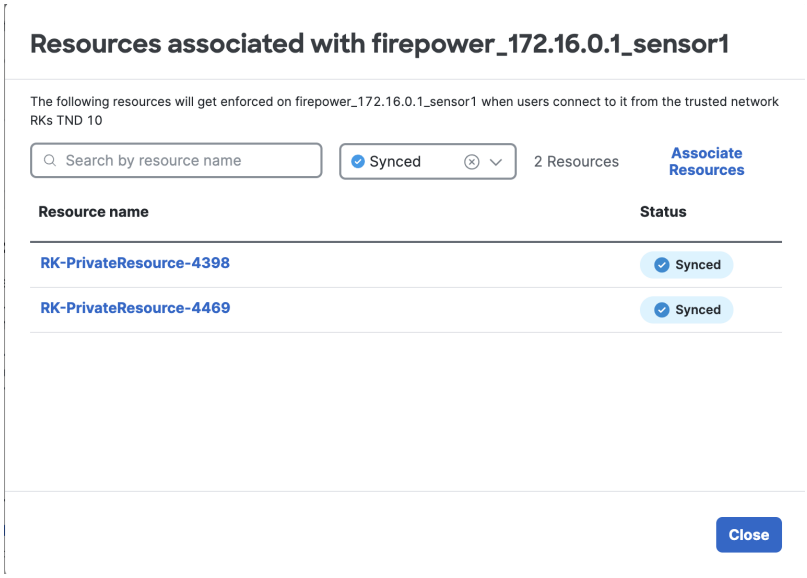
To view a detailed status for each resource and rule associated with a Threat Defense device, complete these actions:

a) Click the numeral in the **Associated Resources** column.

In the slide-in pane, under the **Associated Resources** section, click **View resources associated with this FTD**.

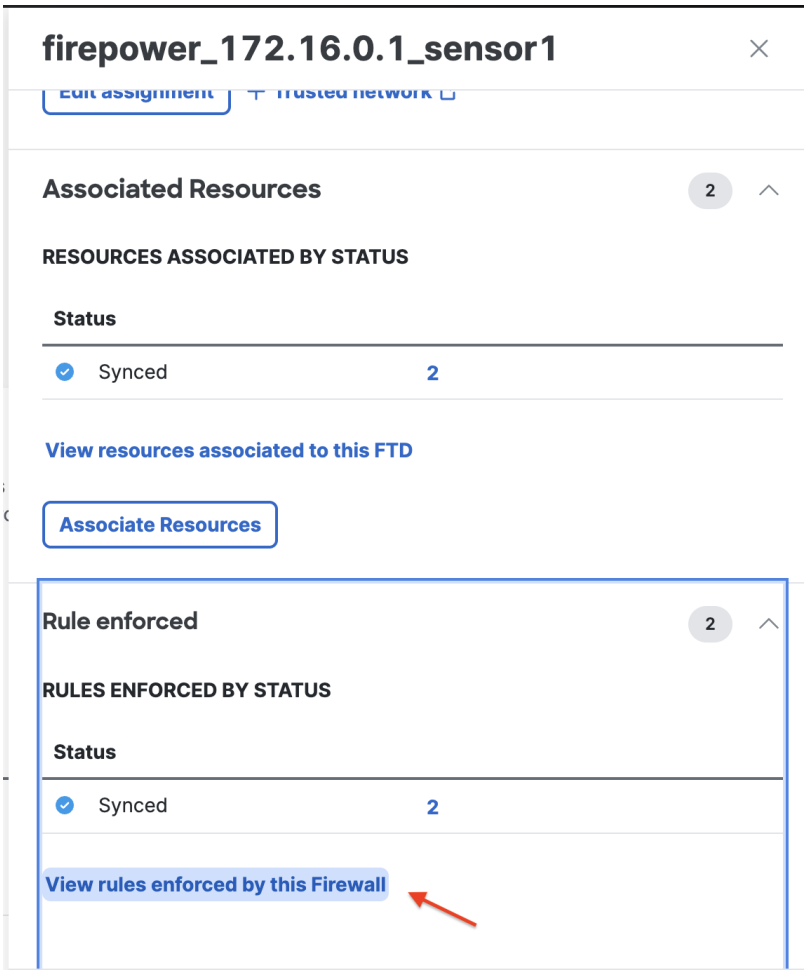


The configuration status of each resource is displayed.

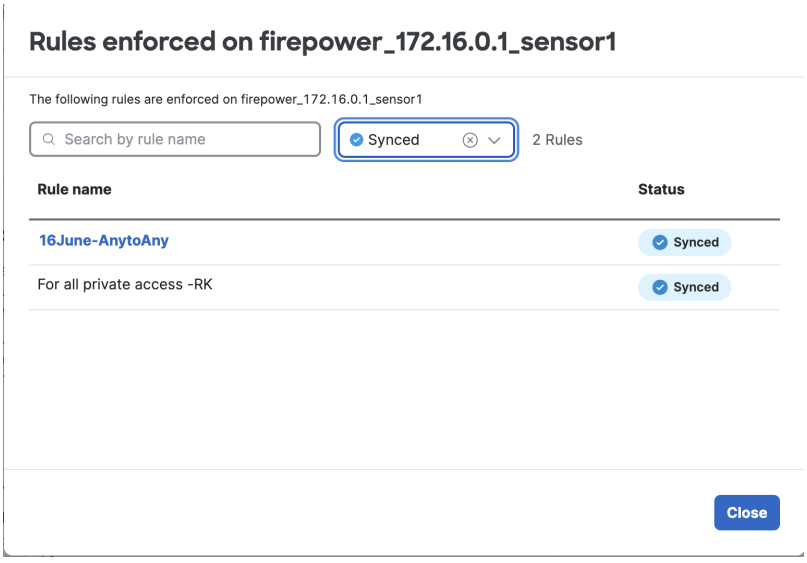


b) To check the configuration status of each rule enforced by the Threat Defense device, click the numeral in the **Rules Enforced** column.

In the slide-in pane, under the **Rules Enforced** section, click **View rules enforced by this Firewall**.



The configuration status of each rule that is enforced is displayed.



Universal ZTNA is now set up, allowing your clients to access private resources in your network securely.

Network connections are configured, and Universal ZTNA is set up to allow secure access to private resources.

Related documentation

This reference provides documentation resources for universal ZTNA configuration components and related products. For a visual walkthrough of the configuration process, watch the videos at this link.

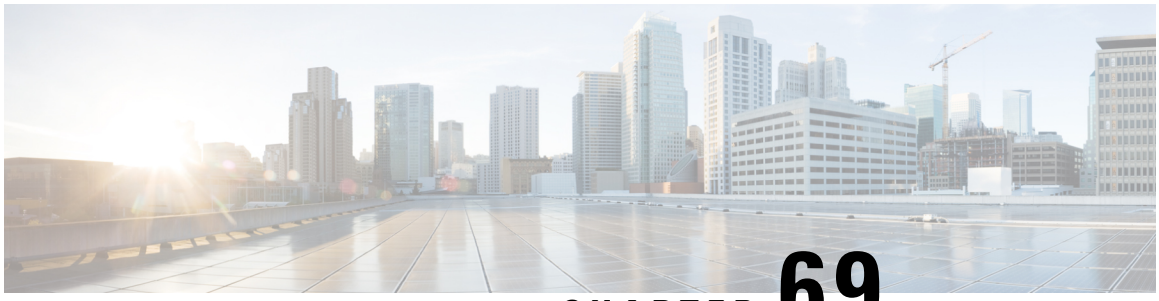
To know more about...	Read this document...
Universal ZTNA Solution	Universal Zero Trust Network Access Solution Guide
Universal ZTNA Configuration	Universal Zero Trust Network Access Configuration Guide
Secure Access	Secure Access Help Center
Security Cloud Control	Security Cloud Control Getting Started Guide
Firewall Threat Defense Health Metrics	Secure Firewall Threat Defense Health Metrics Collected by Firewall Management Center Health Monitor
Firewall Management Center	Secure Firewall Management Center Administration Guide
Cloud-Delivered Firewall Management Center	Managing Threat Defense with Cloud-Delivered Firewall Management Center in Security Cloud Control Release Notes for Cloud-Delivered Management Center
Secure Client	Secure Client Administrator Guide
Supported Device Release	Secure Firewall Threat Defense Release Notes



PART XVII

Access Control Policy Basics

- Access Control Policies, on page 1997
- Security Intelligence, on page 2027
- Access Control Rules, on page 2041
- URL Filtering Rules, on page 2085



CHAPTER 69

Access Control Policies

These topics describe how to work with access control policies:

- Access control policies, on page 1997
- Requirements and prerequisites for access control policies, on page 2002
- Manage access control policies, on page 2002
- History for access control policies, on page 2023

Access control policies

An access control policy is a hierarchical policy that allows you to specify, inspect, and log network traffic. You can use the characteristics of each connection to allow, trust, block, or monitor the connection.

Each managed device can be assigned to one access control policy. The data that the policy's assigned (also known as targeted) devices collect about your network traffic can be used to filter and control that traffic based on:

- simple, easily determined transport and network layer characteristics: source and destination, port, protocol, and so on
- the latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application used, or URL visited
- realm, user, user group, or ISE attribute
- custom Security Group Tag (SGT)
- characteristics of encrypted traffic; you can also decrypt this traffic for further analysis
- whether unencrypted or decrypted traffic contains a prohibited file, detected malware, or intrusion attempt
- time and day (on supported devices)

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance. For example, reputation-based blocking uses simple source and destination data, so it can block prohibited traffic early in the process. In contrast, detecting and blocking intrusions and exploits is a last-line of defense.

Access control policy components

This list describes the main components of an access control policy.

- Name and description: Each access control policy must have a unique name. A description is optional.
- Inheritance settings: Policy inheritance allows you to create a hierarchy of access control policies. A parent (or base) policy defines and enforces default settings for its descendants. A policy's inheritance settings allow you to select its base policy. You can also lock settings in the current policy to force any descendants to inherit them. Descendant policies can override unlocked settings.
- Policy assignment: Every access control policy identifies which devices use it. Each device can use only one access control policy. You can also assign the policy to device templates.
- Rules: Access control rules provide a granular method of handling network traffic. Rules are numbered, starting at 1, including rules inherited from ancestor policies. The system matches traffic to rules from the top down, ascending by rule number.

  Traffic is handled by the first rule for which all conditions match. Conditions can be simple or complex, and often depend on licensing.
- Default action: The default action determines how the system handles and logs traffic that is not handled by any other access control configuration. It can block or trust all traffic without further inspection, or inspect traffic for intrusions and discovery data. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.
- Security Intelligence: This feature acts as the first line of defense against malicious internet content. It enables blocking of connections based on IP address, URL, and domain name reputation intelligence. To ensure continual access to vital resources, you can override block list entries with custom do not block list entries.
- HTTP responses: When a user's website request is blocked, you can display either a generic response page or a custom page. Alternatively, you can warn users but allow them to proceed to the requested site.
- Logging: Logging settings enable configuration of default syslog destinations for the policy. These settings apply to the policy and all included decryption, prefilter, and intrusion policies unless overridden in individual rules or policies.
- Advanced access control options: Advanced policy settings typically require little modification, and defaults are often suitable. Modifiable options include traffic preprocessing, decryption, identity, and various performance configurations.

Access control policy default actions

The default action is the access control policy configuration that determines how the system handles and logs traffic that is not handled by any other access control configuration. The default action can block or trust all traffic without further inspection, or inspect traffic for intrusions and discovery data.

In a simple access control policy, the default action specifies how a device handles all traffic. In a more complex policy, the default action handles traffic that:

- is not fast pathed by the prefilter policy
- is not on a Security Intelligence block list

- is not blocked by the decryption policy (encrypted traffic only)
- matches none of the rules in the policy (except monitor rules, which match and log—but do not handle or inspect—traffic)

Note You cannot perform file or malware inspection on traffic handled by the default action. Logging of connections handled by the default action is initially disabled but can be enabled.

If you are using policy inheritance, the default action for the lowest-level descendant determines how traffic is handled. An access control policy can inherit its default action from the base policy, but this inheritance is not enforced.

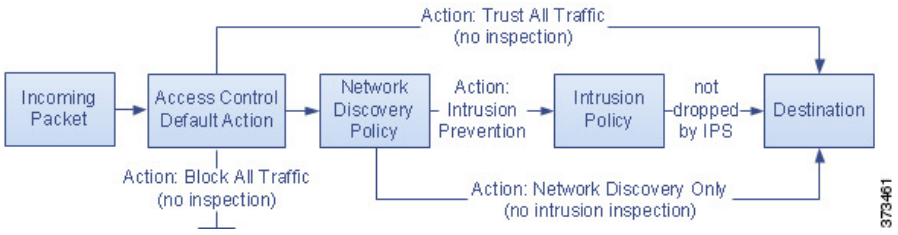
This table describes the types of inspection you can perform on traffic handled by each default action.

Table 115: Types of default actions and inspection options

Default action	Effect on traffic	Inspection type and policy
Access Control: Block All Traffic	Block without further inspection.	None.
Access Control: Trust All Traffic	Trust (allow to its final destination without further inspection).	None.
Intrusion Prevention	Allow, as long as it is passed by the intrusion policy you specify.	Intrusion, using the specified intrusion policy and associated variable set. Discovery, using the network discovery policy.
Network Discovery Only	Allow.	Discovery only, using the network discovery policy.
Inherit from base policy	Defined in base policy.	Defined in base policy.

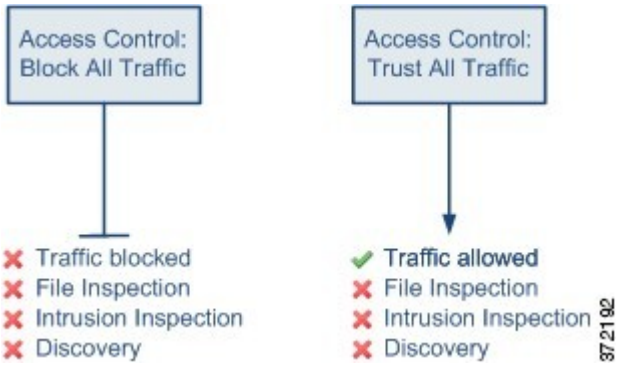
This diagram illustrates the table.

Figure 382: Default action configuration options



This diagram illustrates the **Block All Traffic** and **Trust All Traffic** default actions.

Figure 383: Block all traffic and trust all traffic default actions



This diagram illustrates the **Intrusion Prevention** and **Network Discovery Only** default actions.

Figure 384: Intrusion prevention and network discovery only default actions



Note The purpose of **Network Discovery Only** is to improve performance in a discovery-only deployment. Different configurations can disable discovery if you are only interested in intrusion detection and prevention.

Access control policy inheritance

Access control policy inheritance is a hierarchical policy-based implementation that allows child policies to inherit rules and settings from ancestor or base policies. This structure ensures that mandatory rules defined in parent policies are consistently enforced, while providing the flexibility for child policies to selectively override default configurations.

Access control policy hierarchy structure

You can nest access control policies, where each policy inherits the rules and settings from an ancestor policy. You can enforce this inheritance, or allow lower-level policies to override their ancestors.

Just as you create a domain hierarchy, you can create a corresponding hierarchy of access control policies. A descendant, or child, access control policy inherits rules and settings from its direct parent, or base, policy. That base policy may have its own parent policy from which it inherits rules and settings, and so on.

An access control policy's rules are nested between its parent policy's Mandatory and Default rule sections. This implementation enforces Mandatory rules from ancestor policies, but allows the current policy to write rules that preempt Default rules from ancestor policies.

You can lock these settings to enforce them in all descendant policies:

- Security Intelligence—Connections that are allowed or blocked based on the latest reputation intelligence for IP addresses, URLs, and domain names.
- HTTP response pages—Displaying a custom or system-provided response page when you block a user's website request.
- Advanced settings—Specifying associated subpolicies, network analysis settings, performance settings, and other general options.

When using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

Rule and other policy warnings

Policy and rule editors use icons to mark configurations that could adversely affect traffic analysis and flow. Depending on the issue, the system may warn you when you deploy or prevent you from deploying entirely.

Note Move your pointer over an icon to read the warning, error, or informational text.

Table 116: Policy error icons

Icon	Description	Example
Errors ( )	If a rule or configuration has an error, you cannot deploy until you correct the issue, even if you disable any affected rules.	A rule that performs category and reputation-based URL filtering is valid until you assign a device that does not have a URL Filtering license. At that point, an error icon appears next to the rule, and you cannot deploy until you edit or delete the rule, unassign the device, or enable the license.
Warning ( )	You can deploy a policy that displays rule or other warnings. However, misconfigurations marked with warnings have no effect. If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue.	Preempted rules or rules that cannot match traffic due to misconfiguration have no effect. This includes conditions using empty object groups, application filters that match no applications, excluded LDAP users, invalid ports, and so on. However, if a warning icon marks a licensing error or model mismatch, you cannot deploy until you correct the issue.

Icon	Description	Example
Information ()	Information icons convey helpful information about configurations that may affect the flow of traffic. These issues do not prevent you from deploying.	The system might skip matching the first few packets of a connection against some rules, until the system identifies the application or web traffic in that connection. This allows connections to be established so that applications and HTTP requests can be identified.
Rule Conflict ()	When you enable rule conflict analysis, this icon appears in the rule table for rules that have conflicts.	Conflicts include redundant rules, redundant objects, and shadowed rules. Redundant and shadowed rules do not match traffic because previous rules would already match the criteria. Redundant objects make your rules unnecessarily complex.

Requirements and prerequisites for access control policies

Ensure that your environment meets these requirements before configuring access control policies.

- You must have **Administrator**, **Network Admin**, or **Access Admin** privileges to configure access control policies.
- Custom roles: You can define custom user roles to differentiate between the intrusion configuration in access control policy and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. The existing pre-defined user roles that included the Modify Access Control Policy permission support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions.
  - **Policies** > **Access Control heading** > **Access Control** and choose the **Access Control Policy** > **Modify Access Control Policy** > **Modify Threat Configuration** allows the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for network analysis and intrusion policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions in the policy default action. If a user has this option only, the user can modify no other part of the policy or rule.
  - **Modify Remaining Access Control Policy Configuration** controls the ability to edit all other aspects of the policy.

Manage access control policies

Configure and manage access control policies by editing system-provided policies and creating custom access control policies.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

At the top of the page, there are convenient links to some related features: Object management, Intrusion policies, Network Analysis policies, DNS policies, and policy Import/Export.

Step 2 These actions help you to manage your policy:

- Create—Click **Create Policy**; see Creating a basic access control policy, on page 2004.
- Edit—Click **Edit** ( ); see Editing an access control policy, on page 2004
- Delete—Click **Delete** ( ). You must remove any device assignments before deleting a policy.

  To delete more than one policy at a time, select the check boxes for the policies, then select **Delete Policies** above the table.
- Copy—Select **Clone** from the **More** ( ) menu. Device assignments are not retained in the copy.
- Lock or unlock a policy—See Lock an access control policy, on page 2006.

Step 3 These options help you to analyze, monitor, and customize your policies:

- Analyze policy—Select one or more policy and then click **Analyze Policies** to evaluate access control policies for anomalies such as redundant or shadowed rules, and take action to fix discovered anomalies. The analysis job is sent to the cloud and takes time to complete. See Identify and fix anomalies with policy analyzer & optimizer, on page 2017.

  The Anomalies column shows the results of the analysis. Click the **% Optimizable** link to see the anomalies, or **Re-Analyze** to run the analysis again. The Last Analyzed column shows when Policy Analyzer & Optimizer was last run.

  After completing analysis and optimization, you can download the reports by selecting the following options from the **More** ( ) menu: **Download Last Policy Analysis**, **Remediation History**.

  Note
  To use the policy analysis feature, you must be using a Cloud-Delivered Firewall Management Center or an On-Prem Firewall Management Center connected to Security Cloud Control (Security Cloud Control). If your setup does not meet requirements, the explanatory dialog that opens when you click this button includes an **Integrate** button to help get you started. Policy Analyzer & Optimizer operates in the cloud only.
- Columns—Click the **Show/Hide Columns** icon above the list of rules to select which information to show in the table. Click **Show All/Hide All** to quickly add or remove all listed columns, excepting name and actions. Click **Default** to undo all of your customizations.
- Inheritance—Vertical lines indicate inheritance relationships between policies.
- Report—Select **Generate Report** from the **More** ( ) menu. The report is generated as a background process. Go to the message/notifications center and look in the tasks list. When the report is finished, you can download it from the notification.
- View the audit log—Click **View Audit Log** from the **More** ( ).

Creating a basic access control policy

When you create a new access control policy, it contains default actions and settings. After creating the policy, you are immediately placed in an edit session so that you can adjust the policy to suit your requirements.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Create Policy**.

Step 3 Enter a unique **Name** and, optionally, a **Description**.

Step 4 Optionally, choose a base policy.

If an access control policy is enforced on your domain, this step is not optional. You must choose the enforced policy or one of its descendants as the base policy.

If you select a base policy, the base policy defines the default action and you cannot select a new one in this dialog box. Logging for connections handled by the default action depends on the base policy.

Step 5 When you do not select a base policy, specify the initial **Default Action**:

- **Block all traffic** creates a policy with the **Access Control: Block All Traffic** default action.
- **Intrusion Prevention** creates a policy with the **Intrusion Prevention: Balanced Security and Connectivity** default action, associated with the default intrusion variable set.
- **Network Discovery** creates a policy with the **Network Discovery Only** default action.

When you select a default action, logging of connections handled by the default action is initially disabled. You can enable it later when you edit the policy.

Note
If you want to trust all traffic by default, or if you chose a base policy and do not want to inherit the default action, you can change the default action later.

Step 6 Optionally, choose the devices to assign to the policy. To narrow the devices that appear, enter a search string. The list includes both devices and device templates.

If you want to deploy this policy immediately, you must perform this step.

Step 7 Click **Save**.

The new policy opens for edit. You can add rules to it and make other changes as needed. See Editing an access control policy, on page 2004.

Editing an access control policy

When you edit an access control policy, you should lock it to ensure that your changes do not get overridden by another person who might edit it simultaneously.

You can only edit access control policies that were created in the current domain. Also, you cannot edit settings that are locked by an ancestor access control policy.

To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Note When you edit a policy and do not lock it, a banner message might indicate that other users are currently editing the policy. When one of these users saves changes, you will be prompted to either merge or discard your changes. You should immediately decide what to do. For more information, see Merge concurrent edits to an access control policy, on page 2008 .

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control** .

Step 2 Click **Edit** ( ) next to the access control policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your access control policy.

Note
You can operate on multiple rules at one time by selecting their checkboxes in the left column, then selecting the action you want to perform from the **Select Bulk Rule Actions** drop-down list next to the search box. Bulk editing is available for enabling and disabling, copying, moving, deleting, and editing rules, or viewing hit counts or related events. You can also remove object overlaps in the selected rules.

You can change the following settings or perform these actions:

- Name and description—Click **Edit** ( ) next to the name, make your changes, and click **Save** .
- Default action and settings—Choose a value from the **Default Action** drop-down list, then click **Cog** ( ) , make your changes to the settings, and click **OK** . For detailed information, see Configure the access control default action, on page 2010 .
- Associated policies—To edit or change policies in the packet flow, click the policy type in the packet flow representation below the policy name. You can select the **Prefilter Rules** , **Decryption** , **Security Intelligence** , and **Identity** policies. When necessary, click **Access Control** to return to the access control rules.
- Policy assignment—To identify the managed devices targeted by this policy, or enforce this policy in a subdomain, click the **Assigned: x devices** link. You can assign the policy to devices or device templates.
- Rules—To manage access control rules, and to inspect and block malicious traffic using intrusion and file policies, click **Add Rule** , or right-click an existing rule and select **Edit** or another appropriate action. The actions are also available from the **More** ( ) button for each rule. See Create and edit access control rules, on page 2066 .

- Layout—Use the **Grid/Table View** icon above the list of rules to change the layout. Grid view provides color-coded objects in an easy-to-see layout. Table view provides a summary list so that you can see more rules at once. You can freely switch views without impacting the rules.
- Columns (table view only)—Click the **Show/Hide Columns** icon above the list of rules to select which information to show in the table. Click **Show/Hide Empty Columns** to quickly add, or remove, all columns that have no information, that is, you are not using those conditions in any rule. Click **Revert to Default** to undo all of your customizations.
- Analyze rule logic—You can select the following options from the **Analyze** menu to examine the logic of your rules:
  - **Manage Rule Hit Counts** —To view statistics on how many connections matched each rule. See View rule hit counts, on page 2018 .
  - **Enable/Disable Rule Conflicts** —Select whether you want to see information on whether rules interfere with each other. You can then view the results using the following commands. See Analyzing rule conflicts and warnings, on page 2020
    - **Show Warnings and Errors** —See whether there are rules with configuration issues that you need to address.
    - **Show Policy Warnings** —See whether there are configuration issues with the policy.
    - **Show Rule Conflicts** —See whether you have redundant or shadowed rules. These conflicts could prevent certain rules from ever being matched by connections, meaning either that you need to fix the match criteria, move the rule, or simply delete the rule.

Step 4 To change additional settings for the policy, select one of these options from the **More** drop-down arrow at the end of the packet flow line.

- **Encrypted Visibility Engine** —Whether to enable and configure the encrypted visibility engine. See Configure EVE, on page 2242 .
- **Advanced Settings** —To set preprocessing, decryption, identity, performance, and other advanced options. See Configure advanced settings for the access control policy, on page 2688 .
- **HTTP Responses** —To specify what the user sees in a browser when the system blocks a website request. See Choose HTTP response pages, on page 2105 .
- **Inheritance Settings** —To change the base access control policy for this policy, and to enforce this policy's settings in its descendant policies. See Choose a base access control policy, on page 2011 and Lock settings in descendant access control policies, on page 2012 .
- **Logging** —To set the default logging options for the policy.

Step 5 Click **Save** .

Lock an access control policy

Lock an access control policy to prevent other administrators from editing it, ensuring that your changes will not be invalidated if another administrator edits the policy and saves changes before you save your changes.

If you do not lock a policy, you might have to merge your changes if another user edits the policy simultaneously and saves changes first. For more information, see Merge concurrent edits to an access control policy, on page 2008.

The lock is for the access control policy itself. The lock does not apply to objects used in the policy. For example, another user can edit a network object that is used in a locked access control policy. Your lock remains in place until you explicitly unlock the policy, so you can log out and come back to your edits later.

When locked, other administrators have read-only access to the policy. However, other administrators can assign a locked policy to a managed device.

Before you begin

Any user role that has permission to modify the access control policy has permission to lock it, and to unlock a policy that was locked by another user.

However, the ability to unlock a policy that was locked by another administrator is controlled by the following permission: **Policies** > **Access Control heading** > **Access Control** and choose the **Access Control Policy** > **Modify Access Control Policy** > **Override Access Control Policy Lock**.

If you are using custom roles, your organization might have limited your unlocking abilities by not assigning this permission. Without this permission, only the administrator who locks a policy can unlock it.

Procedure

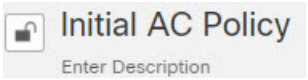
Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to lock or unlock.

In the access control policy list:

- The **Lock Status** column shows whether a policy is already locked, and if so, who locked it. An empty cell indicates that the policy is not locked.
- If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration. Or, it is locked by another user.

Step 3 Click the lock icon next to the policy name to lock or unlock the policy.



If the policy inherits settings from a parent policy, you must choose one of these options when you click the lock icon.

- **Lock/Unlock this policy**—The locking or unlocking is for this policy only.
- **Lock/Unlock all policies in hierarchy**—This policy and all parent policies are locked or unlocked. If a parent policy is already locked by another administrator, you will see a message and you will not be able to lock that parent policy. When unlocking policies, if you have the Override Access Control Policy Lock permission, all parent policies are unlocked even if they were locked by other users.

Merge concurrent edits to an access control policy

Handle conflicts when multiple users make changes to the same access control policy simultaneously, ensuring no critical changes are lost while maintaining policy integrity.

When you allow multiple users to edit an access control policy, you have two choices:

- Have users lock the policy when editing it. This ensures that a single user at a time can make changes. Using locking, you never have to worry about lost or conflicting changes. However, it limits your ability to respond quickly to unexpected events, especially if a user leaves the policy locked while out of the office. For more information about locking policies, see Lock an access control policy, on page 2006.
- Allow users to simultaneously edit the policy. If more than one user has made unsaved edits at a given time, the first user who saves the policy gets all of their changes saved. Other users are immediately notified of the save, and must then merge their changes. The rest of this topic explains this approach.



Note If you are using change management, where edits are done within the scope of a ticket, the ticket locks the policy and concurrent editing is not possible.

Before you begin

To limit the amount of merging you have to do, save the policy early and often.

Procedure

Step 1 Determine if other users are also editing the policy.

When you edit an access control policy, look for this message banner at the top of the page:

```
User name is currently editing this policy.
```

This message indicates that the named user (there might be more than one) currently have the policy open for edit.

Step 2 Watch for notifications that another user saved changes and take action.

This banner message indicates another user saved changes and you need to take immediate action:

```
User name modified this policy and saved changes. Merge Discard.
```

Click one of these links:

- **Merge**—Open a merge window where you can make decisions about which of your changes to keep or to discard, and to identify changes that cannot be merged. See the next step when taking this option.
- **Discard**—Discard all of your changes immediately and start over. If you select this option, the page is refreshed with the latest changes from the other user.

Step 3 Merge changes.

When you click the **Merge** option, a merge window opens with a summary of the total number of observed differences between your edits and the last user's saved edits. No other user's unsaved edits are included. Observations include direct conflicts and informational notifications.

Note
You must complete all merge decisions before you close the window. You cannot wait until later once you click the **Merge** option.

a) Evaluate each conflict and make a decision for each.

When evaluating the change list, consider these points:

- Initially, all observations are shown, but you can select/deselect the filtering options for **Conflict** and **Info** to limit the displayed information. Conflicts arise when users edit the same rule or setting. Informational items are for edits to different rules or settings.
- If there are no conflicts, your cache is immediately updated with the last saved changes and you can proceed. For example, if you are editing rule 1, and the other user saved changes to rule 2, it is unlikely that there are any conflicts.
- For each conflict (other than rule name), the window shows the version of the element saved by the other user (**Version on Firewall Management Center**), and your changes (**Modified Version**). Changes are color-coded for new (e.g. an option that was not previously defined), edited, or deleted policy elements. You must select either to **Discard** your change, or to **Accept Mine** and overwrite the other user's change. When discarding a change, the rule is refreshed with the other user's saved change and yours is removed.
- Rule conflicts are considered at the rule level, not per element within the rule. For example, if you edit destination networks, and the saved changes are for source networks, either your changes or the other user's changes are retained, not both.
- If the conflict involves the name of a rule, the system tries to create a unique name by adding an underscore and number to the name, such as Rule_1. You can alternatively enter a new name in the edit box provided and click **Save**. If you click **Discard**, the generated rule name is used.
- You cannot skip any conflict. For each conflict, you must make a decision to discard or keep your change.
- If the user who saved changes deleted a rule that you edited, the rule is deleted and you do not have the option to keep the rule.
- If the user who saved changes updated a rule, and your change is to delete the rule, the rule is retained: you cannot merge your deletion.
- For these options, if your change conflicts with the saved changes, your change is discarded and merging the change is not an option:
  - Other policy assignments: prefilter, Security Intelligence, identity, decryption
  - Changes to advanced settings, policy logging, EVE, HTTP response pages.
  - Policy default action.
  - Inheritance settings.

b) Click **Close**.

Step 4 Once the merger is complete, you can either save the policy immediately or continue editing it. Saving the policy might reduce the impact of future merges.

Configure the access control default action

Configure the default action for an access control policy when connections do not match specific access control rules.

The default action for an access control policy is applied to any connection that:

- is not fast pathed by the prefilter policy
- is not on a Security Intelligence block list
- is not blocked by the decryption policy (encrypted traffic only)
- matches none of the rules in the policy (except monitor rules, which match and log—but do not handle or inspect—traffic)

Procedure

Step 1 Edit the access control policy whose inheritance settings you want to change; see Editing an access control policy, on page 2004.

Step 2 Select the **Default Action** at the bottom of the rules list.

For detailed information on what each option does, see Access control policy default actions, on page 1998.

Step 3 Click **Cog** (⚙) to configure the default action settings.

You can configure the following options. Click **OK** when finished.

- Logging options—Whether to log the connection. You can **Log at beginning of connection**, **Log at end of connection**, or both. If you select block as the default action, you can log at the beginning of the connection only.
- **Send connection events to**—If you select one of the logging options, select whether to send events to any combination of the following:
  - **Firewall Management Center**—Send events to the manager.
  - **Syslog server**—Send events to the default syslog server configured for the policy. You can configure overrides to specify a different severity level or syslog server destination.

  (Splunk or SIEM syslog server integration) If you have configured Cloud-Delivered Firewall Management Center as the source of connection events in Splunk configuration, choose **Firewall Management Center** under **Send Connection Events to** options. If you have configured Firewall Threat Defense as the connection event source, select at least one syslog destination. For Splunk or any SIEM syslog configuration procedure, see the *Cisco Secure Firewall Management Center Administration Guide*.
- **SNMP trap**—If you enable logging, you can send SNMP traps to an SNMP server. Select an SNMP configuration, or click + to configure a new one.

- **Default action variable set**—If you selected one of the intrusion prevention default actions, select the variable set that should be used with the intrusion policy you selected.

Step 4 Click **Save**.

Manage access control policy inheritance

Inheritance relates to using another policy as a base policy for an access control policy. This allows you to use one policy to define some baseline characteristics that can be applied to multiple policies. To understand how inheritance works, refer to Access control policy inheritance, on page 2000.

Procedure

Step 1 Edit the access control policy whose inheritance settings you want to change; see Editing an access control policy, on page 2004.

Step 2 Manage policy inheritance:

- Change base policy — To change the base access control policy for this policy, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line and proceed as described in Choose a base access control policy, on page 2011.
- Lock settings in descendants — To enforce this policy's settings in its descendant policies, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line and proceed as described in Lock settings in descendant access control policies, on page 2012.
- Inherit settings from base policy — To inherit settings from a base access control policy, proceed as directed in Inherit access control policy settings from the base policy, on page 2012.
- Required in domains — To enforce this policy in a subdomain, click the **Assigned: x devices** link and proceed as described in Enforce an access control policy in a domain, on page 2013.

Choose a base access control policy

Choose a base access control policy to establish policy inheritance, allowing a child policy to inherit rules and settings from its parent while maintaining the ability to modify unlocked configurations.

When you change the base policy for the current access control policy, the system updates the current policy with any locked settings from the new base policy.

Procedure

Step 1 In the access control policy editor, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 Choose a policy from the **Select Base Policy** drop-down list.

To remove inheritance, select **None**.

Step 3 Click **OK**.

Step 4 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Lock settings in descendant access control policies

Lock a setting in an access control policy to enforce the setting in all descendant policies.

Descendant policies can override unlocked settings. When you lock settings, the system saves overrides already made in descendant polices so that the overrides can be restored if you unlock settings again.

Procedure

Step 1 In the access control policy editor, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 In the Child Policy Inheritance Settings area, check the settings you want to lock.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

Step 3 Click **OK** to save the inheritance settings.

Step 4 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Inherit access control policy settings from the base policy

Configure a child access control policy to inherit settings from a base policy or override them with custom configurations.

Child access control policies inherit many settings from their base policy by default. If these settings are unlocked in the base policy, you can override them.

If you later reinherit the settings from the base policy, the Cloud-Delivered Firewall Management Center displays the base policy's settings and dims the controls. However, it saves the overrides you made and restores them if you disable inheritance again.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**, or select **HTTP Responses**, **Logging**, **Encrypted Visibility Engine**, or **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line

Step 2 Check the **Inherit from (*base policy*)** check box for each setting you want to inherit.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

Step 3 Click **Save**.

What to do next

- Deploy configuration changes.

Enforce an access control policy in a domain

Require that every device in a domain use the same base access control policy or one of its descendant policies.

This procedure is relevant in a multi-domain deployment only.

Procedure

Step 1 In the access control policy editor, click the **Assigned: x devices** link.

Step 2 Click **Required on Domains**.

Step 3 Build your domain list:

- Add — Select the domains where you want to enforce the current access control policy, then click **Add** or drag and drop into the list of selected domains.
- Delete — Click **Delete** ( ) next to a leaf domain, or right-click an ancestor domain and choose **Delete Selected**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Step 4 Click **OK** to save the domain enforcement settings.

Step 5 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Assign access control policy to devices

An access control policy specifies the devices that use it. Each device can be assigned to only one access control policy. You can also assign the policy to device templates. Templates are included in the list of available and selected devices.

Procedure

Step 1 In the access control policy editor, click the **Assigned to x devices** link.

Step 2 Build your target list:

- Add — Select one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**.
- Delete — Click **Delete** ( ) next to a single device, or select multiple devices, right-click, then choose **Delete Selection**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Under **Impacted Devices**, the system lists the devices whose assigned access control policies are children of the current policy. Any change to the current policy affects these devices.

Step 3 Assign the access control policy to devices:

- Add — Select one or more **available devices**, then click the **Right Arrow** ( ) icon or drag and drop into the list of **assigned devices**.
- Delete — Click the **Delete** ( ) next to a device, or select multiple devices, and then click the Left Arrow .
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Step 4 Click **OKApply** to save your assigned device settings.

Step 5 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Logging settings for access control policies

Logging settings for access control policies are configuration option that allow you to configure default syslog destinations for the current access control policy. The settings are applicable to the access control policy and all the included decryption, prefilter, and intrusion policies unless the syslog destination settings are explicitly overridden with custom settings in included rules and policies.

To configure logging settings for an access control policy, select **Logging** from the **More** drop-down arrow at the end of the packet flow line.

You can configure default syslog destinations and syslog alert for the access control policy. The settings are applicable to the access control policy and all the included decryption, prefilter, and intrusion policies unless the syslog destination settings are explicitly overridden with custom settings in included rules and policies.

Logging for connections handled by the default action is initially disabled.

Intrusion settings and file and malware settings are effective only after you have selected an option at the top of the page for sending syslog messages generally.

Default syslog settings

Default syslog settings provide two destination options:

- **Destination** > ***Specific Syslog alert object***: Events are sent based on the selected syslog alert as configured using the instructions in *Creating a Syslog alert Response* in the *Cisco Secure Firewall Management Center Administration Guide*. You can select the syslog alert from the list or add one by specifying the name, logging host, port, facility, and severity.

  When using this option, the system sends syslog messages to the server using the Management interface. Ensure there is a route from the Management interface to the syslog server, or messages will not arrive at the server.

- **Destination** > **Use FTD Platform Settings**: Connection or intrusion events are sent with the selected severity to syslog collectors configured in Platform Settings. Using this option, you can unify the syslog configuration by configuring it in platform settings and reusing the settings in access control policy. Severity selected in this section is applied to all connection and intrusion events. The default severity is alert.

Intrusion settings

Intrusion settings include these options:

- **Send syslog messages for intrusion events**: Send intrusion events as syslog messages. The defaults set above are used unless you override them.
- **Show overrides** or **Hide overrides**: If you want to use the default syslog destination and severity, leave these options empty. Otherwise, you can set a different syslog server destination for intrusion events, and change the severity of the events.

File and malware settings

File and malware settings include these options:

- **Send syslog messages for file and malware events**: Send file and malware events as syslog messages. The defaults set above are used unless you override them.
- **Show overrides** or **Hide overrides**: If you want to use the default syslog destination and severity, leave these options empty. Otherwise, you can set a different syslog server destination for file and malware events, and change the severity of the events.

AI defense settings

Associate prefilter, decryption, and identity policies with an access control

Associate subpolicies with an access control policy to create a comprehensive security framework that handles different aspects of network traffic processing.

The easiest way to associate subpolicies to an access control policy is by clicking the policy's link in the packet flow shown at the top of an access control policy. Alternatively, you can use the policy's advanced settings to associate the subpolicy, as described in this topic. The subpolicies that you can associate with access control policy are:

- Prefilter policy—Performs early traffic handling using limited network (layer 4) outer-header criteria.
- Decryption policy—Monitors, decrypts, blocks, or allows application layer protocol traffic encrypted with Secure Socket Layer (SSL) or Transport Layer Security (TLS).
- Identity policy—Performs user identification based on the realm and authentication method associated with the traffic.

Before you begin

Before associating a decryption policy with an access control policy, review the information about TLS server identity discovery in TLS server identity discovery, on page 2692.

Procedure

Step 1 In the access control policy editor, select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 Locate the policy settings area for prefilter, decryption, or identity policies.

Step 3 Click **Edit** ( ) in the corresponding policy settings area.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 Choose a policy from the drop-down list.

If you choose a user-created policy, you can click the edit icon that appears to edit the policy.

Step 5 Click **OK**.

Step 6 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Identify and fix anomalies with policy analyzer & optimizer

Use the Policy Analyzer & Optimizer to evaluate access control policies for anomalies such as redundant or shadowed rules, and take action to fix discovered anomalies.

The Policy Analyzer & Optimizer is hosted in the cloud and is different from the rule analysis available when you are not integrated with the cloud. The non-cloud policy analysis is not available once you integrate with the cloud.

The system automatically performs policy analysis on a daily basis (every 24 hours). You can also manually start an analysis. When you initially enable the service, the system starts an analysis of all existing access control policies.

Note Before optimizing a policy, create a copy of the policy. If you are then dissatisfied with the results of optimization, you can easily reassign the managed devices to the copy and return the system to its starting state.

Before you begin

- If you have enabled change management, Policy Analyzer & Optimizer automatically creates a ticket for the changes, and submits the ticket. The approver must approve the ticket before the changes can be deployed.
- If you are using domains, you cannot click the link in the Anomaly column to see the report. Instead, log into Security Cloud Control and use the feature from that application.
- Policy Analyzer & Optimizer adds rule comments on rules that are updated, disabled, or merged. You can later search on these comments to find optimized rules.
- Changes implemented by Policy Analyzer & Optimizer are reflected in the audit log as API calls under the default name internaladmin.
- After enabling Policy Analyzer & Optimizer, the system evaluates changes every time a rule is added or edited. Users are notified when applying a rule edit if the change will create anomalies such as redundant or shadowed rules.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

If you have already run an analysis, the Anomaly column shows the number of issues with the policy and the percentage the policy can be optimized, and the state of the policy analysis, such as Error or Completed. Last Analyzed shows the date/time when the analysis was run.

Step 2 Select one or more policy, then click **Analyze Policy**.

The analysis runs as a background process in the cloud. When the analysis is complete, the results appear in the Anomaly column.

Notes:

- You can also start an analysis when editing the policy by selecting **Analyze** > **Policy**. Other options from that menu allow you to show hit counts and warnings.
- If you have not connected to the cloud yet, the explanatory dialog that opens when you click this button includes an **Integrate** button to help get you started. Policy Analyzer & Optimizer operates in the cloud only.

Step 3 When the analysis is complete, click the **% Optimizable** link in the Anomaly column to launch Policy Analyzer & Optimizer in the cloud.

When you have done all the actions you want to take, click **Apply Remediations** (in the cloud). You are shown a confirmation of what will be done. Click **Proceed** to implement the changes.

If the initial analysis ended in an error, you could instead click **Re-analyze** to restart the process.

Step 4 Deploy the policy to complete the changes.

If you have change management enabled, the approver must first approved the ticket that contains the remediations before you can deploy them.

View rule hit counts

Hit count indicates the number of times a policy rule or default action has been matched to a connection. The hit count is incremented only for the first packet of a connection that matches a rule. You can use this information to identify the efficacy of your rules. Hit count information is available only for access control and prefilter rules applied to Firewall Threat Defense devices.

Note

- The count persists through reboots and upgrades.
- Counts are maintained by each unit in an HA pair or cluster separately.
- You will not be able to derive the hit count information from a device when deployment or a task is in progress on the device.
- You can also see rule hit count information in the device CLI using the **show rule hits** command.
- If you have accessed the Hit Count page from the Access Control Policy page, you will not be able to view or edit prefilter rules and vice-versa.
- Hit counts are not available for rules that use the Monitor action.

Before you begin

If you use custom user roles, ensure that the roles include these privileges:

- View Device, to see the hit counts.
- Modify Device, to refresh the hit counts.

Follow these steps to view rule hit counts:

Procedure

Step 1 In the access control policy or prefilter policy editor, click **Analyze** > **Manage Rule Hit Counts** on the top-right of the page.

Step 2 On the Hit Count page, select the device from the **Select a device** drop-down list.

If it is not the first time that you are generating hit counts for this device, the last fetched hit count information appears next to the drop-down box. Also, verify the **Last Deployed** time to confirm recent policy changes.

Select **All** to see an aggregate of hit counts across all devices assigned to the policy.

Step 3 If necessary, click **Refresh** ( ) to obtain current hit count data from the selected device.

In the prefilter policy, you might need to click **Fetch Current Hit Count** to get initial hit count data.

You cannot refresh the hit count while deployment to the device is in progress.

Step 4 View and analyze the data.

You can do these actions:

- Click **Prefilter** or **Access Control** to switch between the hit counts for these policies.
- Search for a specific rule by entering a search string in the **Filter** box.
- Broadly limit the list to **Hit Rules** or **Never Hit Rules** by selecting these options in the **Filter by** field. When viewing hit rules, you can further limit the list by selecting a time range in the **In Last** field (for example, in the last 1 day).
- (When viewed from the access control policy.) Clear the hit counts for one or more rule by selecting the checkbox for the rule, then clicking **Clear Hit Counts**. When confirming the action, choose **Clear and Reload** to refresh the hit count data. You can clear hit counts for up to 500 rules at a time. You cannot undo clearing hit counts.

  Note
  Click the checkbox in the table header to select all rules visible in the list. To select a range of rules, select the first rule's checkbox, then **`Shift+click`** the last rule's checkbox; all rules in between are also selected.

- (When viewed from the access control policy.) You can do these actions with individual rules:
  - Edit the rule by clicking **Edit** from the **More** ( ) menu.
  - Delete the rule from the policy by clicking **Delete** from the **More** ( ) menu.
  - Enable or disable the rule by clicking **Enable Rule** or **Disable Rule** from the **More** ( ) menu.
  - Clear the hit count (reset it to zero) for the rule by clicking **Clear Hit Count** from the **More** ( ) menu. You cannot undo this action.
- (When viewed from the access control policy.) You can do these actions with individual rules:
  - Edit the rule by clicking **Edit** ( ).

- Delete the rule from the policy by clicking **Delete** ( ).
- Enable or disable the rule by clicking the **Slider** ( ).
- Clear the hit count (reset it to zero) for the rule by clicking the **X** for the rule. You cannot undo this action.

- Change the displayed columns by clicking and selecting the columns to show. When viewed from the prefilter policy, use the **Cog** ( ) to select the columns to display.
- (When viewed from the prefilter policy.) Click on a rule name to edit it, or click **View** ( ) in the last column to view the rule details. Clicking on the rule name highlights it in the policy page where you can edit it.
- (When viewed from the prefilter policy.) Clear the hit count information (reset it to zero) for a rule by right-clicking the rule and selecting **Clear Hit Count**. You can select multiple rules by using Ctrl+click. You cannot undo this action.

Step 5 Optionally, you can generate a comma-separated values report of the details on the page by clicking **Generate CSV** on the bottom-left of the page.

Step 6 Click **Close** to return to the policy page.

Analyzing rule conflicts and warnings

View warnings and information about rule conflicts to examine the logic of your access control policy and to identify rules that need changes.

Note The feature described in this topic is available only if you do not integrate with the Cisco Security Clound. When you integrate with the cloud, the more powerful Policy Analyzer & Optimizer replaces this feature. See Identify and fix anomalies with policy analyzer & optimizer, on page 2017.

When rules overlap, you can end up with unnecessary rules in the policy, and these rules will never be matched to traffic. The analysis can help you delete unnecessary rules, or identify rules that should be moved or modified so they enforce the desired policy.

Policy warnings and errors point out things that you should understand and perhaps address to ensure that your rules provide the desired services.

Rule conflict analysis identifies these problems:

- Object overlap—An element in a field of a rule is a subset of one or more elements in the same field of the rule. For example, the source field might include a network object for 10.1.1.0/24, and another object for the host 10.1.1.1. Because 10.1.1.1 is within the network covered by 10.1.1.0/24, the object for 10.1.1.1 is redundant and can be deleted, simplifying the rule and saving device memory.
- Redundant rule—Two rules apply the same action to the same type of traffic and removing the base rule would not change the ultimate result. For example, if a rule permitting FTP traffic for a particular network

were followed by a rule allowing IP traffic for that same network, and there were no rules in between denying access, then the first rule is redundant and you can delete it.

- Shadowed rule—This is the reverse of a redundant rule. In this case, one rule will match the same traffic as another rule such that the second rule will never be applied to any traffic because it comes later in the access list. If the action for both rules is the same, you can delete the shadowed rule. If the two rules specify different actions for traffic, you might need to move the shadowed rule or edit one of the two rules to implement your desired policy. For example, the base rule might deny IP traffic, and the shadowed rule might permit FTP traffic, for a given source or destination.

Before you begin

When doing the analysis:

- Only the first conflict for a given rule is identified. Once you fix the problem, the rule might be identified as having a conflict with another rule in the table. However, a rule might have multiple warnings or errors.
- Rule conflict analysis considers source/destination security zone, network, VLAN, and service/port match conditions and action only. It does not consider other match criteria, so an apparently redundant rule might not be completely redundant.
- FQDN network objects cannot be analyzed for conflict, because the IP address of an FQDN cannot be known prior to DNS lookup.
- Disabled rules are ignored.
- Time range attributes are ignored. Rules for different time periods might be marked as redundant when they are not actually redundant for the time ranges.
- Icons for warnings and errors, and rule conflicts when you enable the feature, are shown in the rules table. For a reference to the icons, see Rule and other policy warnings, on page 2001.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control** and edit an access control policy.

Step 2 Do one of these to open the rule conflicts and warnings dialog box:

- To view rule conflicts, click the **Analyze** drop-down and click **Show Rule Conflicts** from the menu.

  You must re-enable rule conflict detection each time you open the policy or make a change and save the policy.

- To view rule warnings and errors, click **Analyze** > **Show Rule Warnings and Errors**.

  After you make a change to the policy, you can refresh the results by clicking **Re-Analyze** next to the **Analyze** button.

- To view policy warnings, click **Analyze** > **Show Policy Warnings**.

Step 3 In the rule conflicts and warnings dialog box:

- Rule warnings and errors are shown on a separate tab from rule conflicts. There is also a separate tab for policy warnings.

- Each tab contains sub-tabs to let you examine individual types of problems, such as redundant vs. shadowed or warnings vs. errors. You can also search for an item.
- **More** (⋮) next to each rule name provides shortcuts to edit, disable, or delete the rule.

Step 4 Click **Close** when finished.

Search for rules

Use search to help you find rules, especially when you have a large number of them.

Procedure

Step 1 When editing an access control policy, build the search string by clicking in the **Search** box.

- For a simple text string search, type the string. The search returns rules that have that string in any column. You cannot do a string search in conjunction with a tag search, such as combining a string search with a source networks search.
- To search on a specific column, start typing the name of the column until the system prompts you with the full name, such as Source Networks, or select the name from the list of searchable fields. When you select the search tag, you can then enter the search string for that tag. For example, **Source Networks 10.1.1.1**.
- When searching for IP addresses within the source/destination columns (but not as a simple text search), you can use the **Seach Inside Subnets and Ranges** option to control whether you eactly match the IP address (unchecked) or match any subnet or range that includes the searched address (checked). For example, with the option checked, searching for 10.1.1.1 will include rules for 10.1.1.0/24.
- Searches in the port field return exact match only.
- When entering mulitple values, separate the values with commas.
- When building a search string with multiple tags, do not include spaces between the tags.
- When you select a tag, you are prompted with values that appear in those columns. Select the values you want to search for.
- You can quickly filter based on some common features by clicking the filter icon to the left of the search box and selecting to show rules with any combination of the following: Allow, Block, Monitor, Intrusion Policy, Time Range, expired rules, , rules that include objects whose definitions overlap..
- To see the rules that apply to a specific device or set of devices, click the filter icon and select the devices. Rules apply to a device if they use security zones that contain at least one interface on the device, or if they do not include security zones.

Step 2 With your cursor at the end of the search string in the search box, press Enter.

The rules that satisfy the search string are highlighted and non-matching rules are hidden. You can deselect **Show Only Matching Rules** to see the entire table, with the highlighted rules within the table. This lets you see the surrounding rules.

Beside the Show Only Matching Rules checkbox is a summary of the total number of rules in the policy compared to the number that match the search string.

Step 3 After your first search, clicking in the search box prompts you with recent searches and tags. You can quickly repeat a search by selecting it, or build similar searches by selecting previous searches or tags and building on them.

Step 4 To close the search and return to the unfiltered and unhighlighted table, click the **X** at the right of the search box. You can also put your cursor at the end of the search string and press the Esc key.

History for access control policies

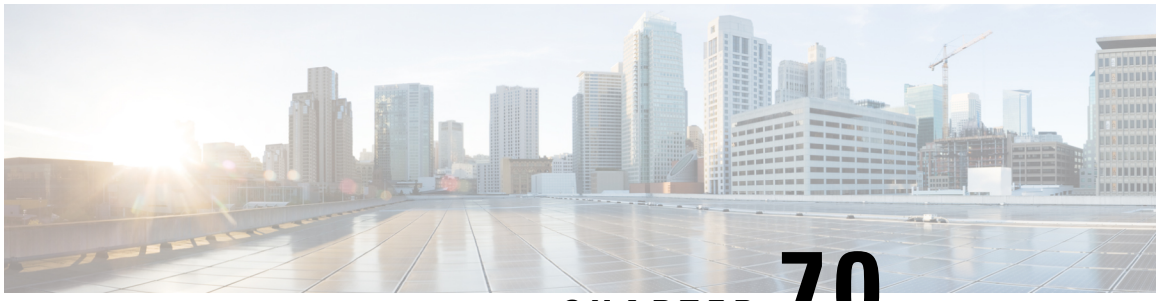
This history table provides a comprehensive history of feature changes for access control policies in Cisco Secure Firewall Management Center across releases.

Table 117: Access control policy feature history by release

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Simultaneous editing of access control policies by multiple users.	10.0.0	Any	In previous releases, if two or more users simultaneously edited an access control policy, the first user who saved would retain their changes, and all other users would immediately lose all of their edits. Now, these users have the ability to selectively merge their changes, and changes that do not conflict with the first user's saved changes will automatically be accepted. This improves collaboration between users and reduces the need to lock the policy during edits.
Object group search performance enhancements.	7.6.0	7.6.0	Object group search is now faster and uses fewer CPU resources. New CLI commands: **clear asp table network-object**, **show asp table network-object**, **debug acl ogs** Modified CLI comments (enhanced output): , **packet-tracer**, **show access-list**, **show object-group** See: Device Settings and *Cisco Secure Firewall Threat Defense Command Reference*

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Granular permissions for modifying access control policies and rules.	7.4.0	Any	You can define custom user roles to differentiate between the intrusion configuration in access control policies and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. When defining user roles, you can select the **Policies** > **Access Control** > **Access Control Policy** > **Modify Access Control Policy** > **Modify Threat Configuration** option to allow the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for Network Analysis and Intrusion Policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions in the policy default action. You can use the **Modify Remaining Access Control Policy Configuration** to control the ability to edit all other aspects of the policy. The existing pre-defined user roles that included the Modify Access Control Policy permission continue to support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions.
New access control policy user interface and rule conflict analysis.	7.3.0	Any	The access control policy user interface introduced in 7.2 is now the default interface. You can also enable rule conflict analysis to help identify redundant rules and objects, and shadowed rules that cannot be matched due to previous rules in the policy.
Access control policy locking.	7.2.0	Any	You can lock an access control policy to prevent other administrators from editing it. Locking the policy ensures that your changes will not be invalidated if another administrator edits the policy and saves changes before you save your changes. Any user who has permission to modify the access control policy has permission to lock it. We added an icon to lock or unlock a policy next to the policy name while editing the policy. In addition, there is a new permission to allow users to unlock policies locked by other administrators: Override Access Control Policy Lock. This permission is enabled by default in the Administrator, Access Admin, and Network Admin roles.
Rule hit counts persist over reboot.	7.2.0	Any	Rebooting a managed device no longer resets access control rule hit counts to zero. Hit counts are reset only if you actively clear the counters. In addition, counts are maintained by each unit in an HA pair or cluster separately. You can use the **show rule hits** command to see cumulative counters across the HA pair or cluster, or see the counts per node. We modified the device CLI command: **show rule hits**.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Usability improvements for the access control policy.	7.2.0	Any	There is a new user interface available for the access control policy. You can continue to use the legacy user interface, or you can try out the new user interface. The new interface has both a table and a grid view for the rules list, the ability to show or hide columns, enhanced search, infinite scroll, a clearer view of the packet flow related to policies associated with the access control policy, and a simplified add/edit dialog box for creating rules. You can freely switch back and forth between the legacy and new user interfaces while editing an access control policy.
DNS filtering	7.0.0 6.7.0 (experimental)	Any	If URL filtering is enabled and configured, a new option to enhance category and reputation filtering efficacy is enabled by default for each new access control policy. For more information, see DNS filtering, on page 2099 and subtopics. The Advanced tab of access control policy has a new option under General Settings: **Enable reputation enforcement on DNS traffic**.
TLS server identity discovery	6.7.0	Any	Enable access control policies to evaluate URL and application conditions when a client connects to a TLS 1.3-enabled server. TLS server identity discovery enables these conditions to be evaluated without decrypting traffic. Enabling this feature can impact device performance, depending on model. The Advanced tab page of access control policy has new options: • Warning is displayed on the Advanced tab; moving the slider to the right enables TLS server identity discovery. • New option on the Advanced tab page: **TLS Server Identity Discovery**.
New Security Intelligence categories	—	Any	These categories were introduced at about the time of the 6.6 release, but are not specific to 6.6: • banking_fraud • high_risk • ioc • link_sharing • malicious • newly_seen • spyware



CHAPTER 70

Security Intelligence

The following topics provide an overview of Security Intelligence, including use of lists for blocking and allowing traffic and basic configuration.

- Security intelligence, on page 2027
- License requirements for security intelligence, on page 2028
- Requirements and prerequisites for security intelligence, on page 2028
- Best practices for security intelligence, on page 2028
- Guidelines and limitations for security intelligence, on page 2029
- Security intelligence sources, on page 2029
- Configure security intelligence, on page 2030
- Security intelligence monitoring, on page 2037
- Override security intelligence blocking, on page 2038
- Troubleshooting security intelligence, on page 2039

Security intelligence

Security Intelligence is an early line of defense mechanism that

- uses reputation intelligence to quickly block connections to or from IP addresses, URLs, and domain names
- operates as the first phase of access control, before the system performs more resource-intensive evaluation, and
- improves performance by quickly excluding traffic that does not require inspection through Security Intelligence block listing.

Security intelligence features and limitations

Security Intelligence provides access to regularly updated intelligence feeds from Cisco. Sites representing security threats such as malware, spam, botnets, and phishing appear and disappear faster than you can update and deploy custom configurations.

You can refine Security Intelligence Block listing with Do Not Block lists and monitor-only Block lists. These mechanisms exempt traffic from being blocked by a Block list, but do **not** automatically trust or fastpath matching traffic. Traffic added to a Do Not Block list or monitored at the Security Intelligence stage is intentionally subject to further analysis with the rest of access control.

Note You cannot use a Block list to block fastpathed traffic. Prefilter evaluation occurs before Security Intelligence filtering. Fastpathed traffic bypasses all further evaluation, including Security Intelligence.

Related Topics

Security intelligence, on page 3021

License requirements for security intelligence

Threat Defense License

IPS

Classic License

Protection

Requirements and prerequisites for security intelligence

Model support

Any

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Important You must apply the Network Discovery policy on the device for a successful application of the SI policy.

Best practices for security intelligence

Configure your access control policies to block threats detected by Cisco-provided Security Intelligence feeds. See Configuration example: Security intelligence blocking, on page 2036.

If you want to supplement the Cisco-provided Security Intelligence feeds with custom threat data, or manually block emerging threats:

- For IP addresses, use custom Security Intelligence lists and feeds, or Network objects. To create these, see Security intelligence, on page 3021 and Networks, on page 2982, and their subtopics. To use them for Security Intelligence, see Configure security intelligence, on page 2030. Network objects used in Security Intelligence policy require an IPS license.
- For URLs and domains, use custom Security Intelligence lists and feeds, *not* objects or groups. See details at Manual URL filtering options, on page 2102.
- You can also add entries to a Block list from events. See Global and domain security intelligence lists, on page 3023.

To test new feeds, or for passive deployments, set the action from block to monitor only. See Security intelligence monitoring, on page 2037.

If you need to exclude specific sites or addresses from Security Intelligence blocking, see Override security intelligence blocking, on page 2038.

If your Firepower deployment is integrated with Cisco XDR and you use custom Security Intelligence lists and feeds, be sure to update Security Services Exchange with these lists and feeds. For details, see instructions for configuring auto-promotion of events in the Security Services Exchange online help.

System-provided Security Intelligence categories may change over time and without notification; you should plan to check periodically for changes, and modify your policies accordingly.

You should also configure URL filtering, a separate feature with separate licensing requirements, for further protection against malicious sites. See URL Filtering Rules, on page 2085.

To reduce latency and CPU overhead, avoid referencing a single, massive SI feed object across your entire security policy. Instead, utilize customized objects to create smaller, targeted subsets of your intelligence data.

Using URL/Domain feeds instead of IP blocklists can reduce CPU overhead because domain filtering is generally more efficient and less resource-intensive than managing massive IP blocklists.

Excluding trusted IPs by adding known good or essential partner IP addresses to an "Do Not Block" list in Security Intelligence (SI) feeds helps reduce CPU overhead as the firewall or security device bypasses unnecessary scanning and deeper inspection for these IPs.

Guidelines and limitations for security intelligence

Be aware of these limitations when configuring and using Security Intelligence:

- The Security Intelligence configuration that includes detection signatures or malware URLs might lead to false threat alerts when the Firewall Threat Defense troubleshoot files are scanned.
- If a TCP connection attempt is initiated from an IP address that is blocked due to its poor reputation, a TCP reset (RST) packet is sent. This causes the connection to be terminated instantly, preventing the establishment of the TCP session and signalling the rejection to the sender.

Security intelligence sources

Security Intelligence sources are data providers that

- supply regularly updated intelligence feeds for domains, URLs and IP addresses

- enable custom blocking and allow lists for specific sites or addresses, and
- provide global lists for immediate event-based blocking and allowing.

Types of security intelligence sources

Security Intelligence uses these source types:

- **System-provided feeds**: Cisco provides access to regularly updated intelligence feeds for domains, URLs and IP addresses. For more information, see Security intelligence, on page 3021.
- **Third-party feeds**: Optionally, supplement Cisco-provided feeds with third-party reputation feeds, which are dynamic lists that the Cloud-Delivered Firewall Management Center downloads from the internet on a regular basis. See Custom security intelligence feeds, on page 3029.
- **Custom Block lists or feeds (or objects)**: Block specific IP addresses, URLs, or domain names using a manually-created list or feed (for IP addresses, you can also use network objects.)

  For example, if you become aware of malicious sites or addresses that are not yet blocked by a feed, add these sites to a custom Security Intelligence list and add this custom list to the Block list in the Security Intelligence tab of your access control policy, as described in Custom Security Intelligence lists, on page 3031 and Configure security intelligence, on page 2030.

  For IP addresses, you can optionally use network objects rather than lists or feeds for this purpose; for information, see Networks, on page 2982. (For URLs, using lists and feeds is strongly recommended over other methods.)
- **Custom Do Not Block lists or feeds**: Override Security Intelligence blocking for specific sites or addresses. See Override security intelligence blocking, on page 2038.
- **Global Block lists (one each for Network, URL and DNS)**: While reviewing events, you can immediately add an event's IP address, URL, or domain to the applicable Global Block List so that Security Intelligence will handle future traffic from that source. See Global and domain security intelligence lists, on page 3023.
- **Global Do Not Block lists (one each for Network, URL and DNS)**: While reviewing events, you can immediately add an event's IP address, URL, or domain to the applicable Global Do Not Block List if you do not want Security Intelligence to block future traffic from that source. See Global and domain security intelligence lists, on page 3023.

Configure security intelligence

Configure Security Intelligence filtering to automatically block or monitor traffic based on threat intelligence feeds and custom block or allow lists. This provides an additional layer of security by filtering known malicious IP addresses, URLs, and domains before they reach your network resources.

Each access control policy has Security Intelligence options. You can add network objects, URL objects and lists, and Security Intelligence feeds and lists to a Block list or Do Not Block list, and constrain any of these by security zone. You can also associate a DNS policy with your access control policy, and add domain names to a Block or Do Not Block list.

The number of objects in the Do Not Block lists plus the number in the Block lists cannot exceed 125 network objects, or 32767 URL objects and lists.

Before you begin

- Tip: For guidance on minimum configuration recommendations, see also Configuration example: Security intelligence blocking, on page 2036.
- To ensure that ALL options are available to select, add at least one managed device to your management center.
- In passive deployments, or if you want to set Security Intelligence filtering to monitor-only, enable logging
- Configure a DNS policy to take Security Intelligence action for domains. For more information, see DNS Policies for Security Intelligence, on page 2557.

Follow these steps to configure Security Intelligence filtering in an access control policy:

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Choose the type of objects to configure.

- Click **Networks** to add network objects (IP addresses).

  Note
  Network objects used in a Security Intelligence policy require a IPS license.

- Click **URLs** to add URL objects.

Step 3 Find the **Available Objects** you want to add to the Block or Do Not Block list.

- Search the available objects by typing in the **Search by name or value** field. Clear the search string by clicking **Reload** ( ) or **Clear** ( ).
- If no existing list or feed meets your needs, click **Add** ( ), select **New Network List** or **New URL List**, and proceed as described in Create custom security intelligence feeds, on page 3030 or Upload new custom security intelligence lists to the Cloud-Delivered Firewall Management Center, on page 3032.
- If no existing object meets your needs, click **Add** ( ), select **New Network Object** or **New URL Object**, and proceed as described in Create a network object, on page 2985.

Security Intelligence ignores IP address blocks using a `/0` netmask.

Step 4 Choose one or more **Available Objects** to add.

Step 5 (Optional) Choose an **Available Zone** to constrain the selected objects by zone.

You cannot constrain system-provided Security Intelligence lists by zone.

Note

The **Any** zone for an SI list applies only to interfaces that are part of a security zone. However, an exception is that if a device does not have any interfaces associated with a security zone, then the **Any** zone will match any interface.

For example, if you have five interfaces on a device and none of them are associated with a security zone, any SI list that is assigned to the **Any** zone will be inspected against traffic on ALL interfaces on the device. If you add one interface to a security zone on that device, it effectively would remove SI inspection on the other four interfaces, where the zone is set to **Any** for an SI list. If you add the other four interfaces to a security zone, they will be evaluated by the SI list attached to the **Any** zone.

Step 6 Click **Add to Do Not Block list** or **Add to Block list**, or click and drag the selected objects to either list.

To remove an object from a Block or Do Not Block list, click **Delete** ( ) To remove multiple objects, choose the objects and right-click to **Delete Selected**.

Step 7 (Optional) Set objects on the Block list to monitor-only by right-clicking the object under **Block List**, then choosing **Monitor-only (do not block)**.

You cannot set system-provided global Security Intelligence lists to monitor only.

Step 8 Choose a DNS policy from the **DNS Policy** drop-down list.

Step 9 Click **Save**.

Security Intelligence filtering is now configured for your access control policy. The system will automatically block or monitor traffic based on the configured Block and Do Not Block lists, providing enhanced protection against known threats.

What to do next

- Deploy configuration changes.

Related Topics

Security intelligence, on page 3021
Snort restart scenarios, on page 194

Security intelligence options

Use the Security Intelligence tab in the access control policy editor to configure network (IP address) and URL Security Intelligence, and to associate the access control policy with a DNS policy in which you have configured Security Intelligence for domains.

Available objects

Available objects include:

- Security Intelligence categories populated by the system-provided feed.

  For details, see Security intelligence categories, on page 2034.

- System-provided Global Block and Do Not Block lists.

  For descriptions, see Security intelligence sources, on page 2029.

- Security Intelligence lists and feeds that you create under Object > Object Management > Security Intelligence.

  For descriptions, see Security intelligence sources, on page 2029.

- Network and URL objects that are configured on the respective pages under Object > Object Management. These are different from the Security Intelligence objects in the previous bullet.

  For details about network objects, see Networks, on page 2982. (For URLs, use Security Intelligence lists or feeds rather than objects.)

  Note that range and group network object types are not supported in Security Intelligence policy.

Available zones

Except for the system-provided Global lists, you can constrain Security Intelligence filtering by zone.

For example, to improve performance, you might want to target enforcement to specific zones rather than all interfaces. As a more specific example, you can block spam only for a security zone that handles email traffic.

To enforce Security Intelligence filtering for an object on multiple zones, you must add the object to the Block or Do Not Block list separately for each zone.

Note that range and group network object types are not supported in Security Intelligence policy.

DNS policy

In order to match DNS traffic using Security Intelligence, you must select a DNS policy for your Security Intelligence configuration.

Using Block or Do Not Block lists, or monitoring traffic based on a DNS list or feed, also requires that you:

- Configure DNS Security Intelligence lists and feeds. See Security intelligence, on page 3021.
- Create a DNS policy. See Create basic DNS policies, on page 2561 for more information.
- Configure DNS rules that reference your DNS lists or feeds. See Create and edit DNS rules, on page 2563 for more information.
- Because you deploy the DNS policy as part of your access control policy, you must associate both policies. See Requirement: deploy DNS policy configuration changes, on page 2575 for more information.

Do not block list

See Override security intelligence blocking, on page 2038.

To select all objects in the list, right-click an object.

Block list

See Configuration example: Security intelligence blocking, on page 2036 and other topics in this chapter.

For explanations of the visual indicators in the Block list, see Block list icons, on page 2036.

To select all objects in the list, right-click an object.

Logging

Security Intelligence logging, enabled by default, logs all blocked and monitored connections handled by an access control policy's assigned devices. However, the system does not log Do Not Block list matches; logging of connections on the Do Not Block list depends on their eventual disposition. Logging must be enabled for connections on the Block list before you can set objects on that list to monitor-only.

To enable, disable, or view logging settings, right-click an object in the Block list.

Related Topics

Global and domain security intelligence lists, on page 3023
Security intelligence lists and multitenancy, on page 3024

Security intelligence categories

Security Intelligence categories are determined by the system-provided feeds described in Security intelligence, on page 3021.

These categories are used in the following locations:

- The Networks sub-tab on the Security Intelligence tab of an access control policy
- The URLs sub-tab beside the Networks tab on the Security Intelligence tab of an access control policy
- In a DNS policy on the DNS tab in the DNS rule configuration page
- In events generated when traffic matches Block or Monitor configurations in the above locations

Note If your organization is using Secure Firewall Threat Intelligence Director: When viewing events, you may see categories that indicate that the action was taken by TID, such as `TID URL Block`.

Categories are updated by Talos from the cloud, and this list may change independently of Firepower releases.

Table 118:

Security Intelligence Category	Description
Attackers	Active scanners and hosts known for outbound malicious activity
Banking_fraud	Sites that engage in fraudulent activities that relate to electronic banking
Bogon	Bogon networks and unallocated IP addresses
Bots	Sites that host binary malware droppers
CnC	Sites that host command-and-control servers for botnets
Cryptomining	Hosts providing remote access to pools and wallets for the purpose of mining cryptocurrency

Security Intelligence Category	Description
Dga	Malware algorithms used to generate a large number of domain names acting as rendezvous points with their command-and-control servers
Exploitkit	Software kits designed to identify software vulnerabilities in clients
High_risk	Domains and hostnames that match against the OpenDNS predictive security algorithms from security graph
Ioc	Hosts that have been observed to engage in Indicators of Compromise (IOC)
Link_sharing	Websites that share copyrighted files without permission
Malicious	Sites exhibiting malicious behavior that do not necessarily fit into another, more granular, threat category
Malware	Sites that host malware binaries or exploit kits
Newly_seen	Domains that have recently been registered, or not yet seen via telemetry. **Attention** Currently, this category does not have any active feed and is reserved for future use.
Open_proxy	Open proxies that allow anonymous web browsing
Open_relay	Open mail relays that are known to be used for spam
Phishing	Sites that host phishing pages
Response	IP addresses and URLs that are actively participating in malicious or suspicious activity
Spam	Mail hosts that are known for sending spam
Spyware	Sites that are known to contain, serve, or support spyware and adware activities
Suspicious	Files that appear to be suspicious and have characteristics that resemble known malware
Tor_exit_node	Hosts known to offer exit node services for the Tor Anonymizer network

Block list icons

The visual indicators may appear in the Block list on the Security Intelligence tab in an access control policy.

Icon or Visual Indicator	Description
Block ( )	The object is set to block.
Monitor ( )	The object is set to monitor-only. See Security intelligence monitoring, on page 2037
An object is displayed in strikethrough text	The same object is also on the Do Not Block list, which overrides the block.

Configuration example: Security intelligence blocking

Configure Security Intelligence blocking to protect your network from known malicious IP addresses, URLs, and domains using regularly updated threat intelligence feeds.

Configure your access control policy to block all threats detectable by the system's regularly updated Security Intelligence feeds.

The number of objects in the Block lists plus the number in the Do Not Block lists cannot exceed 125 network objects, or 32767 URL objects and lists.

Before you begin

- To ensure that all options are available to select, add at least one managed device to your management center.
- Configure a DNS policy to block all Security Intelligence threat categories for domains. For more information, see DNS Policies for Security Intelligence, on page 2557.
- If you have, or will have, custom lists of entities to block, create a Security Intelligence object of each type (URLs, DNS, Networks.) See Security intelligence, on page 3021.

Procedure

Step 1 Click **Policies** > **Access Control heading** > **Access Control**.

Step 2 Create a new access control policy or edit an existing policy.

Step 3 In the access control policy editor, click **Security Intelligence**.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 Click **Networks** to add blocking criteria for IP addresses.

a) Scroll down in the Networks list and select all of the threat categories listed below the Global lists.
b) If applicable, select the security zones for which you want to block these threats.
c) Click **Add to Block List**.

d) If you have created custom lists or feeds with addresses to block, add those to the Block List using the same steps as above.

Step 5 Click **URLs** to add blocking criteria for URLs, and repeat the steps you followed for Networks.

Step 6 Choose a DNS policy from the **DNS Policy** drop-down list; see DNS policies, on page 2557.

Step 7 Click **Save**.

Your access control policy now blocks malicious IP addresses, URLs, and domains based on Security Intelligence threat feeds, providing enhanced protection against known threats.

What to do next

- Enable logging for these connections
- Deploy configuration changes.
- For additional protection, configure URL filtering to block malicious URLs. See URL Filtering Rules, on page 2085.

Security intelligence monitoring

Security Intelligence monitoring is a security feature that

- logs connection events for traffic that would have been blocked by Security Intelligence, but does not block the traffic,
- enables testing of feeds before implementation, and
- optimizes performance in passive deployments.

Use cases for security intelligence monitoring

Monitoring is especially useful for:

- Testing feeds before you implement them.

  Consider a scenario where you want to test a third-party feed before you implement blocking using that feed. When you set the feed to monitor-only, the system allows connections that would have been blocked to be further analyzed by the system, but also logs a record of each of those connections for your evaluation.

- Passive deployments, to optimize performance.

  Managed devices that are deployed passively cannot affect traffic flow; there is no advantage to configuring the system to block traffic. Additionally, because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

Note If configured, Secure Firewall Threat Intelligence Director may impact the action taken (Monitor or Block.)

To configure Security Intelligence monitoring:

After you configure Security Intelligence blocking following the instructions in Configuration example: Security intelligence blocking, on page 2036, right-click each applicable object in the Block list and choose **Monitor-only**. You cannot set system-provided Security Intelligence lists to monitor only.

Override security intelligence blocking

Use Do Not Block lists to exempt specific domains, URLs, or IP addresses from being blocked by Security Intelligence lists or feeds.

You can use Do Not Block lists for various purposes:

- Override the occasional false-positive block in a reputable Security Intelligence feed
- Inspect specific traffic in depth instead of blocking it early based on reputation
- Exempt otherwise-restricted transactions based on zone from Security Intelligence blocking

  For example, you can add an improperly classified URL to a Do Not Block list, but then restrict the Do Not Block list object using a security zone used by those in your organization who need to access those URLs. That way, only those with a business need can access the URLs on the Do Not Block list.



Note Entries on a Do Not Block list are simply exemptions from the block list. Any connection that passes the Security Intelligence policy is subject to the access control rules. Thus, an entry in the Do Not Block list can subsequently be blocked by an access control rule or intrusion policy. Your Do Not Block entries should always be exemptions from your block lists.

Procedure

Step 1 Option 1: Add an IP address, URL, or domain from an event to the Global Do Not Block List. See Global and domain security intelligence lists, on page 3023.

Step 2 Option 2: Use a custom Security Intelligence list or feed.

a) Create the custom Security Intelligence list or feed. See Custom Security Intelligence lists, on page 3031 or Create custom security intelligence feeds, on page 3030.

b) For IP addresses (Networks) and URLs: Edit your access control policy, click the Security Intelligence tab, then click the custom list or feed in the Networks or URLs sub-tab, then click **Add to Do Not Block List**.

c) Save your changes.

d) For domains (DNS): See the "DNS Policy" section in the Security intelligence options, on page 2032 topic.

e) Deploy your changes.

The specified domains, URLs, or IP addresses are now exempted from Security Intelligence blocking and will be subject to access control rules instead.

Troubleshooting security intelligence

See the following topics for troubleshooting Security Intelligence.

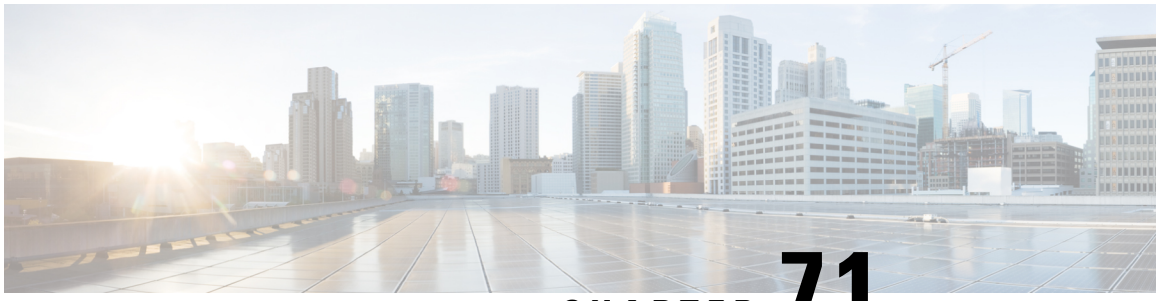
Security intelligence categories missing from the Available Options List

Security Intelligence categories (such as CnC or Exploitkit) display in the Networks tab under Available Options only after you add at least one managed device to your management center.

You must add a device in order to pull all TALOS feeds that provide the Security Intelligence categories.

If the expected category does not appear in Security Intelligence, verify that it is not a URL filtering category, which uses a different set of categories.

To see URL filtering categories, look at the **URLs** tab in an access control rule.



CHAPTER 71

Access Control Rules

The following topics describe how to configure access control rules.

- Access control rules, on page 2041
- Requirements and prerequisites for access control rules, on page 2052
- Guidelines and limitations for access control rules, on page 2053
- Best practices for application control, on page 2054
- Best practices for access control rules, on page 2058
- Access control rule management, on page 2062
- Examples for access control rules, on page 2080

Access control rules

Within an access control policy, access control rules provide a granular method that

- handles network traffic across multiple managed devices within an access control policy
- provides conditions that match traffic in a specified order, and
- determines actions to monitor, trust, block, or allow matching traffic.

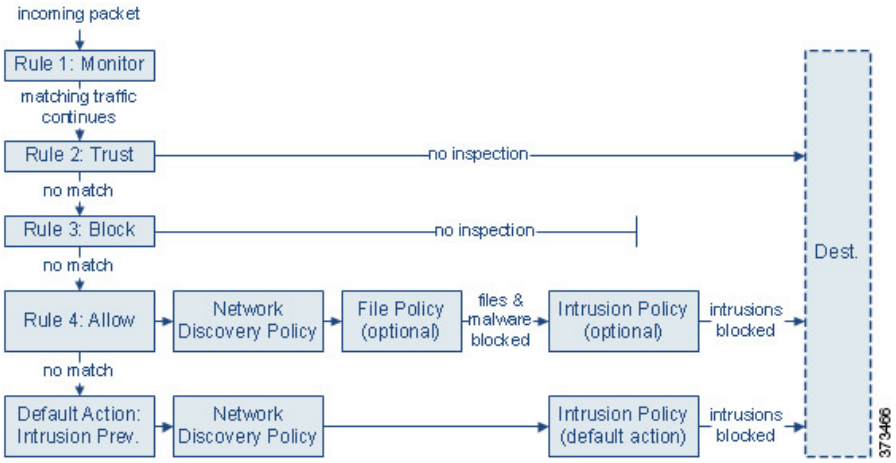
Rule evaluation and traffic handling

The system matches traffic to access control rules in the order you specify. In most cases, the system handles network traffic according to the first access control rule where all the rule's conditions match the traffic.

Each rule also has an action, which determines whether you monitor, trust, block, or allow matching traffic. When you allow traffic, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

Note Security Intelligence filtering, decryption, user identification, and some decoding and preprocessing occur before access control rules evaluate network traffic.

This scenario summarizes the ways that traffic can be evaluated by access control rules in an inline, intrusion prevention deployment.



In this scenario, traffic is evaluated as follows:

- **Rule 1: Monitor** evaluates traffic first. Monitor rules track and log network traffic. The system continues to match traffic against additional rules to determine whether to permit or deny it. (However, see an important exception and caveat at Access control rule monitor action, on page 2046.)
- **Rule 2: Trust** evaluates traffic next. Matching traffic is allowed to pass to its destination without further inspection, though it is still subject to identity requirements and rate limiting. Traffic that does not match continues to the next rule.
- **Rule 3: Block** evaluates traffic third. Matching traffic is blocked without further inspection. Traffic that does not match continues to the final rule.
- **Rule 4: Allow** is the final rule. For this rule, matching traffic is allowed; however, prohibited files, malware, intrusions, and exploits within that traffic are detected and blocked. Remaining non-prohibited, non-malicious traffic is allowed to its destination, though it is still subject to identity requirements and rate limiting. You can configure Allow rules that perform only file inspection, or only intrusion inspection, or neither.
- **Default Action** handles all traffic that does not match any of the rules. In this scenario, the default action performs intrusion prevention before allowing non-malicious traffic to pass. In a different deployment, you might have a default action that trusts or blocks all traffic, without further inspection. (You cannot perform file or malware inspection on traffic handled by the default action.)

Traffic you allow, whether with an access control rule or the default action, is automatically eligible for inspection for host, application, and user data by the network discovery policy. You do not explicitly enable discovery, although you can enhance or disable it. However, allowing traffic does not automatically guarantee discovery data collection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy; additionally, application discovery is limited for encrypted sessions.

Note that access control rules handle encrypted traffic when your decryption configuration allows it to pass, or if you do not configure decryption. However, some access control rule conditions require unencrypted traffic, so encrypted traffic may match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Access control rule management

The rules table of the access control policy editor allows you to add, edit, categorize, search, filter, move, enable, disable, delete, and otherwise manage access control rules in the current policy.

Properly creating and ordering access control rules is a complex task, but one that is essential to building an effective deployment. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations.

Use the search bar to filter the list of access control policy rules. You can deselect the **Show Only Matching Rules** option to see all rules. Matched rules are highlighted.

For each access control rule, the policy editor displays its name, a summary of its conditions, the rule action, and icons that communicate the rule's inspection options or status.

The system provides specific icons to communicate the inspection options or status of each rule:

- **Time Range Option** ( )
- **Intrusion policy** ( )
- **File policy** ( )
- **Logging** ( )
- **Warning** ( )
- **Errors** ( )
- **Rule Conflict** ( )

Disabled rules are dimmed and marked (disabled) after the rule name.

To create or edit a rule, refer to Create and edit access control rules, on page 2066 .

Tip A right-click menu provides shortcuts to many rule management options, including editing, deleting, moving, enabling, and disabling.

Related Topics

Access control rule components, on page 2043
Best practices for access control rules, on page 2058

Access control rule components

In addition to its unique name, each access control rule has basic components that determine how the system processes and handles network traffic.

State

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Position

Rules in an access control policy are numbered, starting at 1. If you are using policy inheritance, rule 1 is the first rule in the outermost policy. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Rules can also belong to a section and a category, which are organizational only and do not affect rule position. Rule position goes across sections and categories.

Section and category

To help you organize access control rules, every access control policy has two system-provided rule sections, Mandatory and Default. To further organize access control rules, you can create custom rule categories inside the Mandatory and Default sections.

If you are using policy inheritance, the current policy's rules are nested between its parent policy's Mandatory and Default sections.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can be simple or complex; their use often depends on license.

Traffic must meet all of the conditions specified in a rule. For example, if the Application condition specifies HTTP but not HTTPS, the URL category and reputation conditions will not apply to HTTPS traffic.

Applicable time

You can set rules to operate during specified days and times

Action

A rule's action determines how the system handles matching traffic. You can monitor, trust, block, or allow (with or without further inspection) matching traffic. The system does not perform deep inspection on trusted, blocked, or encrypted traffic.

Inspection

Deep inspection options govern how the system inspects and blocks malicious traffic you would otherwise allow. When you allow traffic with a rule, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

Logging

Logging settings govern records that the system keeps for handled traffic. You can log sessions at connection start, end, or both; log connections to databases, syslog, or SNMP trap servers. You can log connections to the database, as well as to a syslog or an SNMP trap server.

Comments

Each time you save changes to an access control rule, you can add comments.

Related Topics

Best practices for access control rules, on page 2058

Access control rule management, on page 2043
Create and edit access control rules, on page 2066
Access control rule actions, on page 2046
Access control rule conditions, on page 2068
Deep inspection using file and intrusion policies, on page 2049
Access control rule comments

Access control rule order

Access control rule order is a traffic processing method that

- numbers rules in an access control policy starting at 1
- matches traffic to access control rules in top-down order by ascending rule number, and
- handles network traffic according to the first access control rule where all the rule's conditions match the traffic.

In most cases, except for Monitor rules, the system does not continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule.

To help you organize access control rules, every access control policy has two system-provided rule sections, Mandatory and Default. To further organize, you can create custom rule categories inside the Mandatory or Default sections. After you create a category, you cannot move it, although you can delete it, rename it, and move rules into, out of, within, and around it. The system assigns rule numbers across sections and categories.

If you use policy inheritance, the current policy's rules are nested between its parent policy's Mandatory and Default rule sections. Rule 1 is the first rule in the outermost policy, not the current policy, and the system assigns rule numbers across policies, sections, and categories.

Any predefined user role that allows you to modify access control policies also allows you to move and modify access control rules within and among rules categories. You can, however, create custom roles that restrict users from moving and modifying rules. Any user who is allowed to modify access control policies can add rules to custom categories and modify rules in them without restriction.



Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

Note Proper access control rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs. For more information, refer to Best practices for ordering rules, on page 2059.

Access control rule actions

An access control rule action is a configuration setting that determines how the system handles and logs matching traffic, with options to monitor, trust, block, or allow traffic with or without further inspection.

Default action behavior

The access control policy's *default action* handles traffic that does not meet the conditions of any access control rule with an action other than Monitor.

Access control rule monitor action

The **Monitor** action is an access control rule action that forces connection logging regardless of how matching traffic is eventually handled. It does not affect traffic flow by permitting or denying traffic immediately, and allows traffic to be matched against additional rules to determine the final action.

When a connection matches a monitor rule, the next non-monitor rule that the connection matches determines traffic handling and any further inspection. If there are no additional matching rules, the Firewall Threat Defense device uses the default action.

There is an exception: If a monitor rule contains layer 7 conditions, such as an application condition, the system allows early packets to pass and the connection to be established (or the SSL handshake to complete), even if a subsequent rule would block the connection. Early packets are not evaluated against subsequent rules, so to prevent uninspected traffic, you can specify an intrusion policy for these packets in the access control policy's advanced settings. For more information, refer to Inspection of packets that pass before traffic is identified, on page 2918. After the system completes its layer 7 identification, it applies the appropriate action to the remaining session traffic.

Access control rule trust action

The **Trust** action is the access control rule action that allows traffic to pass without deep inspection or network discovery, still subjects trusted traffic to identity requirements and rate limiting, and bypasses additional inspection services while maintaining basic security controls.

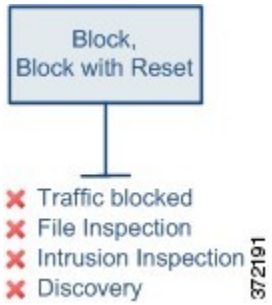


Note

- Some protocols, such as FTP and SIP, use secondary channels, which the system opens through the process of inspection. In some cases, trusted traffic can bypass all inspection, and these secondary channels cannot be opened properly. If you run into this problem, change the trust rule to **Allow**.
- For trust rules with logging options disabled, end-of-flow events are still generated in the system. However, the events are not visible on the event pages.
- Because access control rules are evaluated after other policies, such as decryption, trusting a connection does not necessarily mean it will be fast-pathed with no inspection. For example, if a connection matches both a decryption rule that requires decryption, and a trust access control rule, the connection is decrypted and inspected as needed, prior to being allowed by the trust rule. Trust simply means no additional inspection, such as intrusion inspection, will be applied. If your intention is to allow a connection uninspected, either use the prefilter policy to fast path the connection, or ensure that no other policy applies inspection services to the connection.

Access control rule blocking actions

A block action is an access control rule action that denies network traffic without further inspection and prevents the use of file, intrusion, or network discovery policies for blocked traffic.



Block and **Block with reset** actions immediately stop traffic from entering the network. Block with reset rules reset the connection except for web requests met with an HTTP response page. The HTTP response page is used to notify users when their web request is blocked, but it cannot be displayed if the connection is immediately reset.

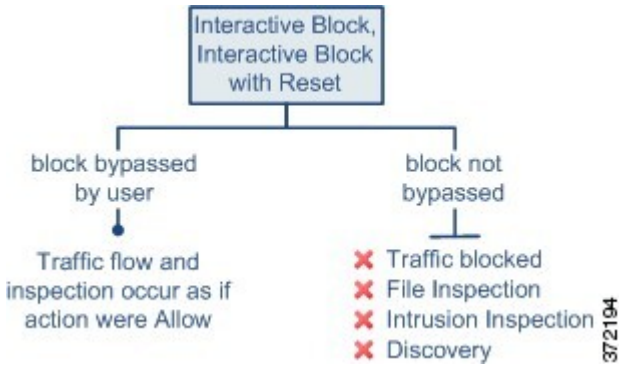
For more information, refer to HTTP response pages, on page 2103.

Related Topics

HTTP response pages, on page 2103

Access control rule interactive blocking actions

The **Interactive Block** and **Interactive Block with reset** actions are access control rule actions that give web users a choice to continue to their intended destinations.



If a user bypasses the block, the rule mimics an allow rule. Therefore, you can associate interactive block rules with file and intrusion policies, and matching traffic is also eligible for network discovery.

If a user does not (or cannot) bypass the block, the rule mimics a block rule. Matching traffic is denied without further inspection.

Note If you enable interactive blocking, you cannot reset all blocked connections. This is because the response page cannot display if the connection is immediately reset. Use the **Interactive Block with reset** action to (non-interactively) block-with-reset all non-web traffic, while still enabling interactive blocking for web requests.

For more information, refer to HTTP response pages, on page 2103.

Related Topics

Decryption rule blocking actions, on page 2222

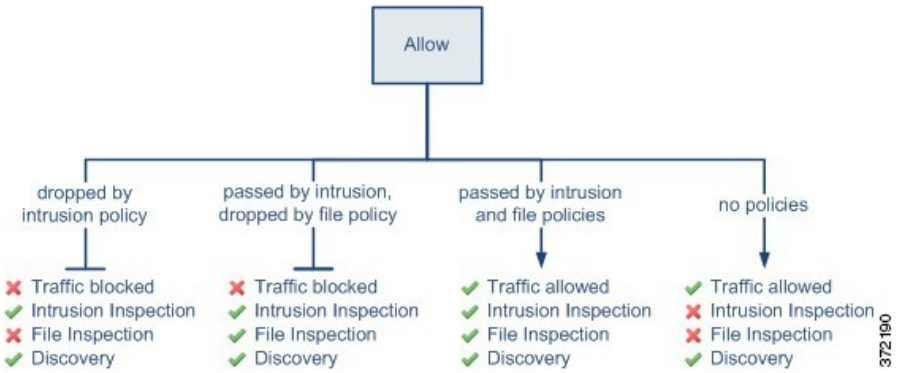
Access control rule allow action

An access control rule allow action is a rule action that allows matching traffic to pass. However, traffic is still subject to identity requirements and rate limiting.

Optionally, you can use deep inspection to further inspect and block unencrypted or decrypted traffic before it reaches its destination:

- You can use an intrusion policy to analyze network traffic according to intrusion detection and prevention configurations, and drop offending packets depending on the configuration.
- You can perform file control using a file policy. File control allows you to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols.
- You can perform network-based advanced malware protection (AMP), also using a file policy. malware defense can inspect files for malware, and block detected malware depending on the configuration.

This diagram illustrates the types of inspection performed on traffic that meets the conditions of an Allow rule (or a user-bypassed Interactive Block rule). Notice that file inspection occurs before intrusion inspection; blocked files are not inspected for intrusion-related exploits.



For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with an access control rule. You can, however, configure one without the other. Without a file policy, traffic flow is determined by the intrusion policy; without an intrusion policy, traffic flow is determined by the file policy.

Regardless of whether the traffic is inspected or dropped by an intrusion or file policy, the system can inspect it using network discovery. However, permitting traffic does not guarantee that discovery inspection occurs. The system only performs discovery for connections involving IP addresses explicitly monitored by your network discovery policy. Also, application discovery is limited for encrypted sessions.

Deep inspection using file and intrusion policies

Deep inspection is a security mechanism that uses intrusion and file policies as the last line of defense before traffic is allowed to its destination.

Policy types and capabilities

Deep inspection uses two types of policies to analyze traffic:

- *Intrusion policies* govern the system's intrusion prevention capabilities.
- *File policies* govern the system's file control and malware defense capabilities.

Access control occurs before deep inspection; access control rules and the access control default action determine which traffic is inspected by intrusion and file policies.

By associating an intrusion or file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy, a file policy, or both.

In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.

Note By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

To associate intrusion and file policies with an access control rule, see:

- Access Control Rules: Intrusion Policy Selection
- Configure an access control rule to perform malware protection, on page 2592

For more information about intrusion policies, refer to Network Analysis and Intrusion Policy Basics.

For more information about file policies, refer to File Policies for Network Malware Protection, on page 2583.

File and intrusion inspection order

File and intrusion inspection order defines the traffic processing sequence that performs file inspection before intrusion inspection for any single connection handled by an access control rule, preventing files blocked by a file policy from being inspected for intrusions while prioritizing simple file blocking by type over malware inspection.

For a connection matching an Allow or Interactive Block rule, the system determines traffic flow based on the following configurations:

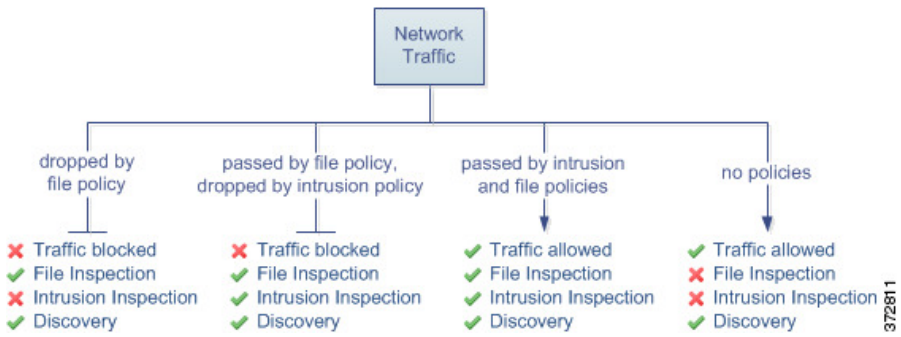
- Without a file policy, traffic flow is determined by the intrusion policy.
- Without an intrusion policy, traffic flow is determined by the file policy.
- Without either, allowed traffic is inspected by network discovery only.

In your access control policy, you can associate multiple Allow and Interactive Block rules with different intrusion and file policies to match inspection profiles to various types of traffic.

Note The system does not perform any kind of inspection on trusted traffic. Although configuring an Allow rule with neither an intrusion nor file policy passes traffic like a Trust rule, Allow rules let you perform discovery on matching traffic.

The diagram below illustrates the types of inspection you can perform on traffic that meets the conditions of either an Allow or user-bypassed Interactive Block access control rule. For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with a single access control rule.

Figure 385: Traffic inspection flowchart



For any single connection handled by an access control rule, file inspection occurs before intrusion inspection. That is, the system does not inspect files blocked by a file policy for intrusions. Within file inspection, simple blocking by type takes precedence over malware inspection and blocking.

For example, consider a scenario where you normally want to allow certain network traffic as defined in an access control rule. However, as a precaution, you want to block the download of executable files, examine downloaded PDFs for malware and block any instances you find, and perform intrusion inspection on the traffic.

You create an access control policy with a rule that matches the characteristics of the traffic you want to provisionally allow, and associate it with both an intrusion policy and a file policy. The file policy blocks the download of all executables, and also inspects and blocks PDFs containing malware:

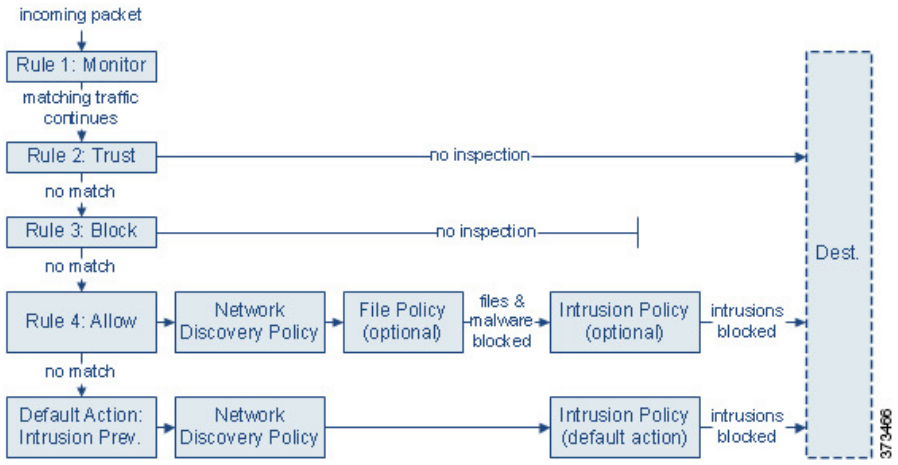
- First, the system blocks the download of all executables, based on simple type matching specified in the file policy. Because they are immediately blocked, these files are subject to neither malware nor intrusion inspection.
- Next, the system performs malware cloud lookups for PDFs downloaded to a host on your network. Any PDFs with a malware disposition are blocked, and are not subject to intrusion inspection.
- Finally, the system uses the intrusion policy associated with the access control rule to inspect any remaining traffic, including files not blocked by the file policy.

Note Until a file is detected and blocked in a session, packets from the session might be subject to intrusion inspection.

How access control traffic handling works with intrusion and file policies

Summary

This diagram shows the flow of traffic in an inline intrusion prevention and malware defense deployment, as governed by an access control policy that contains four different types of access control rules and a default action.



Workflow

In the scenario above, the first three access control rules in the policy—Monitor, Trust, and Block—cannot inspect matching traffic. Monitor rules track and log but do not inspect network traffic, so the system continues to match traffic against additional rules to determine whether to permit or deny it. (However, see an important

exception and caveat at Access control rule monitor action, on page 2046.) Trust and Block rules handle matching traffic without further inspection of any kind, while traffic that does not match continues to the next access control rule.

The fourth and final rule in the policy, an Allow rule, invokes various other policies to inspect and handle matching traffic, in the following order:

1. Discovery: Network Discovery Policy—First, the network discovery policy inspects traffic for discovery data. Discovery is passive analysis and does not affect the flow of traffic. Although you do not explicitly enable discovery, you can enhance or disable it. However, allowing traffic does not automatically guarantee discovery data collection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy.
2. malware defense and File Control: File Policy—After traffic is inspected by discovery, the system can inspect it for prohibited files and malware. malware defense detects and optionally blocks malware in many types of files, including PDFs, Microsoft Office documents, and others. If your organization wants to block not only the transmission of malware files, but all files of a specific type (regardless of whether the files contain malware), file control allows you to monitor network traffic for transmissions of specific file types, then either block or allow the file.
3. Intrusion Prevention: Intrusion Policy—After file inspection, the system can inspect traffic for intrusions and exploits. An intrusion policy examines decoded packets for attacks based on patterns, and can block or alter malicious traffic. Intrusion policies are paired with variable sets, which allow you to use named values to accurately reflect your network environment.
4. Destination—Traffic that passes all the checks described above passes to its destination.

An Interactive Block rule (not shown in the diagram) has the same inspection options as an Allow rule. This is so you can inspect traffic for malicious content when a user bypasses a blocked website by clicking through a warning page.

Traffic that does not match any access control rules in the policy with an action other than Monitor is handled by the default action. In this scenario, the default action is an Intrusion Prevention action, which allows traffic to its final destination as long as it is passed by the intrusion policy you specify. In a different deployment, you might have a default action that trusts or blocks all traffic without further inspection. Note that the system can inspect traffic allowed by the default action for discovery data and intrusions, but not prohibited files or malware. You cannot associate a file policy with the access control default action.

Note Sometimes, when a connection is analyzed by an access control policy, the system must process the first few packets in that connection, allowing them to pass, before it can decide which access control rule (if any) will handle the traffic. However, so these packets do not reach their destination uninspected, you can specify an intrusion policy (in the advanced settings for the access control policy) to inspect these packets and generate intrusion events.

Requirements and prerequisites for access control rules

Ensure that your deployment meets these requirements before you configure access control rules.

Model support

Any

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin
- You can define custom user roles to differentiate between the intrusion configuration in access control policy and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. The existing pre-defined user roles that included the Modify Access Control Policy permission support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions. The granular permissions are:
  - **Policies** > **Access Control heading** > **Access Control** and choose the **Access Control Policy** > **Modify Access Control Policy** > **Modify Threat Configuration** allows the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for Network Analysis and Intrusion Policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions in the policy default action. If a user has this option only, the user can modify no other part of the policy or rule.
  - **Modify Remaining Access Control Policy Configuration** controls the ability to edit all other aspects of the policy.

Guidelines and limitations for access control rules

Consider these limitations when you configure rules in the access control policy:

- There is a limit of 1000 rules that can be shown at a time on the current page. Thus, if you have a very large number of rules, such as 3000 rules within a single category, actions like selecting all rules in a category and deleting them does not delete all of the rules. You might need to select/delete rules again to delete all the rules that you are targeting.
- If you edit an access control rule that is actively in use, the changes do not apply to established connections at deploy-time. The updated rule is used to match against future connections. However, if the system is actively inspecting a connection (for example, with an intrusion policy), it will apply changed matching or action criteria to existing connections.

  For Firewall Threat Defense, you can ensure that your changes apply to all current connections by using the Firewall Threat Defense **clear conn** CLI command to end established connections. Note that you should only do this if it is acceptable to end those connections, on the assumption that the sources for the connections will then attempt to reestablish the connection and thus be matched appropriately against the new rule.

- VLAN tags in access rules only apply to inline sets; they cannot be used in access rules applied to firewall interfaces.
- To use fully-qualified domain name (FQDN) network objects as source or destination criteria, you must also configure DNS for the data interfaces in the platform settings policy. The system does not use the management DNS server setting to do lookups for FQDN objects used in access control rules.

  Note that controlling access by FQDN is a best-effort mechanism. Consider the following points:

  - Whenever possible, use Security Intelligence or URL filtering instead of FQDN rules.
  - Because DNS replies can be spoofed, only use fully trusted DNS servers.
  - Some FQDNs, especially for very popular servers, can have hundreds if not thousands of IP addresses, and these can frequently change. Because the system uses cached DNS lookup results, users might get addresses that are not yet in the cache, and their connections will not match the FQDN rule. Rules that use FQDN network objects function effectively only for names that resolve to fewer than 100 addresses.

    We recommend that you do not create network object rules for an FQDN that resolves to more than 100 addresses, as the likelihood of the address in a connection being one that has been resolved and available in the DNS cache on the device is low. For these cases, use a URL-based rule instead of an FQDN network object rule.
  - For popular FQDNs, different DNS servers can return a different set of IP addresses. Thus, if your users use a different DNS server than the one you configure, FQDN-based access control rules might not apply to all IP addresses for the site that are used by your clients, and you will not get the intended results for your rules.
  - Some FQDN DNS entries have very short time to live (TTL) values. This can result in frequent recompilation of the lookup table, which can impact overall system performance.
  - If more than 8 FQDNs resolve to the same IP address, the system cannot reliably match traffic to the rules for those FQDN. At most 8 FQDNs per IP address can be handled.
- The maximum objects per match criteria per access control rule is 200. For example, you can have up to 200 network objects in a single access control rule.
- Block with reset can be CPU-intensive for rules that drop a high rate of TCP traffic. When the system drops TCP traffic with this action, it also generates TCP reset packets, which can affect device performance during high-volume drops. For access control rules that are expected to drop large amounts of TCP traffic, use **Block** instead of **Block with reset**, unless TCP reset behavior is required.

  If TCP reset behavior is required for denied inbound TCP traffic, you can also evaluate using a regular **Block** rule with the **service resetinbound** command through FlexConfig for the applicable interface.

Best practices for application control

These topics discuss our recommended best practices for controlling applications with access control rules.

Recommendations for application control

Ensure that adaptive profiling is enabled

If adaptive profiling is not enabled (its default state), access control rules cannot perform application control.

Enable application detectors automatically

If no detector is enabled for an application you want to detect, the system automatically enables all system-provided detectors for the application. If none exist, the system enables the most recently modified user-defined detector for the application.

Configure your policy to examine the packets that must pass before an application is identified

The system cannot perform application control before *both* of these conditions occur:

- A monitored connection is established between a client and server
- The system identifies the application in the session

This identification should occur in 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted. If you configure the access control rule to use application default ports, the application rule can be enforced without allowing initial packets to pass.

If early traffic matches all other criteria but application identification is incomplete, the system allows the packet to pass and the connection to be established (or the SSL handshake to complete). After the system completes its identification, the system applies the appropriate action to the remaining session traffic.

To ensure that your system examines these initial packets, select an intrusion policy in the **Intrusion Policy used before Access Control rule is determined** options in the access control policy advanced settings.

Understand limitations to identifying applications

A server must adhere to the protocol requirements of an application for the system to be able to recognize it. For example, if you have a server that sends a keep-alive packet rather than an ACK when an ACK is expected, that application might not be identified, and the connection will not match the application-based rule. Instead, it will be handled by another matching rule or the default action. This might mean that connections you want to allow can be denied instead. If you run into this problem, and you cannot fix the server to follow the protocol standards, you need to write a non-application-based rule to cover traffic for that server, for example, by matching the IP address and port number.

Create separate rules for URL and application filtering

Create separate rules for URL and application filtering whenever possible, because combining application and URL criteria can lead to unexpected results, especially for encrypted traffic.

Rules that include both application and URL criteria should come after application-only or URL-only rules, unless the application+URL rule is acting as an exception to a more general application-only or URL-only rule.

Place URL rules before application and other rules

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of these criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.

Handling application traffic packets without payloads

When performing access control, the system applies the default policy action to packets that do not have a payload in a connection where an application is identified.

Handling referred application traffic

To handle traffic referred by a web server, such as advertisement traffic, match the referred application rather than the referring application.

Controlling application traffic that uses multiple protocols

Some applications use multiple protocols. To control their traffic, make sure your access control policy covers all relevant options. These applications typically include the application name in each aspect that you can control. Also see Application-specific requirements for filtering, on page 2057.

Controlling evasive application traffic

See Application-specific requirements for filtering, on page 2057.

Best practice for choosing between application matching and port matching

In traditional firewalls, you can match traffic based on OSI layers 3 (protocol) and 4 (transport), such as IP or TCP/80. All traffic on a given port (or a whole protocol) is then either allowed or blocked based on that rule's action.

Application criteria, on the other hand, are OSI layer 7. Different applications can use the same TCP/UDP port. By using application criteria, you can selectively allow or block different applications on the same port without allowing or blocking all applications on that port.

Whether you use port-based or application-based criteria in your rule can impact rule performance. Because TCP/UDP ports are quickly identifiable in packets, the system can match the correct rule on the first packet. With application-layer criteria, it can take 3-5 packets to identify the specific application (if you do not also specify ports).

Consider these best practices:

- If you want to handle all of the traffic on a given TCP/UDP port in the same way for the specified interfaces and networks, use port-based matching. For example, to handle all SSH traffic the same way, select the SSH port (TCP/22) on the Ports tab.
- If you want to narrowly match a specific application that uses the same port as other applications, select the application on the Applications tab. This is how you can handle web applications, which all use TCP/80 or 443, so that you can selectively block or allow certain web applications without blocking or allowing all web applications.
- If you want to control application use by a user group, select the user group on the Users tab and the application on the Application tab. For example, you could block the Gaming applications category for members of the Contractor user group.
- For rules that go from a less secure network (such as the internet) to a more secure network (such as an internal protected network), use the Ports tab whenever possible. For example, you can allow/block ICMP traffic from the internet to an internal network.
- Avoid mixing Port tab and Applications tab specifications in a single rule. Instead, when you select an application, leave the Port tab empty, and instead specify the port as **Application Default** on the Applications tab. This will target the rule to the right ports (per application) and enhance how quickly the right application can be matched to the rule.

- If you do have a mix of port-based and application-based rules, place the port-based rules higher in the rules list, so that connections can be matched to those rules first. Protocol and port can be more quickly identified than applications.

Application-specific requirements for filtering

Office 365 admin portal

If the access policy has logging enabled at the beginning as well as at the end of a connection, the first packet will be detected as Office 365 and the end of connection will be detected as Office 365 Admin Portal. This should not affect blocking.

Skype

To control Skype traffic, choose the **Skype** tag from the **Application Filters** list rather than selecting individual applications. This ensures that the system can detect and control all Skype traffic the same way.

GoToMeeting

In order to fully detect GoToMeeting, your rule must include all of these applications:

- GoToMeeting
- Citrix Online
- Citrix GoToMeeting Platform
- LogMeIn
- STUN

Zoho mail

To control Zoho mail, choose *both* **Zoho** and **Zoho mail** from the Available Application list.

Evasive application detection considerations

For evasive applications such as Bittorrent, Tor, Psiphon, and Ultrasurf, only the highest-confidence scenarios are detected by default. If you need to take action on this traffic (such as block or implement QoS), it may be necessary to configure more aggressive detection with better effectiveness. To do this, contact TAC to review your configurations as these changes may result in false positives.

WeChat

It is not possible to selectively block WeChat Media if you allow WeChat.

RDP file transfer

If allowing the RDP (Remote Desktop Protocol) application does not allow file transfers, ensure that the rule for RDP includes both the TCP and UDP port 3389. RDP file transfer uses UDP.

Best practices for access control rules

Configuring and ordering access control rules properly is essential to protect your network. Consider these best practices to achieve optimal results:

- Always evaluate device compatibility of access control rules before deployment.
- Ensure resources, such as physical memory and processors, can handle expanded rule criteria.

Note Ensure the system evaluates all rules together and creates a set of criteria for devices when you deploy configuration changes. You cannot deploy to a device if these criteria exceed its resources.

General best practice for access control

General best practices

- Use a prefilter policy to provide early blocking for unwanted traffic, and to fastpath traffic that does not benefit from access control inspection. For more information, refer to Best practices for fastpath prefiltering, on page 2542.
- Access control rules are deployed as access control entries (ACE) on the device. To minimize the number of access control entries created per access control rule, and improve overall performance, enable object group search for each device. Object group search is a device setting, not an access control policy setting, so you must edit each device to ensure the feature is enabled. For more information, refer to Configure object group search, on page 1190.
- Whenever possible, combine multiple network objects into a single object group. The system automatically creates an object group (during deployment) when you select more than one object (for source or destination separately). Selecting existing groups can avoid object group duplication and reduce the potential impact on CPU usage when there are a large number of duplicate objects.
- Certain features, including URL filtering, application detection, and rate limiting, must allow some packets to pass in order for the system to identify the traffic.
- If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.

Licensing and feature support

- Although you can configure the system without licensing your deployment, many features require that you enable the appropriate licenses before you deploy.
- Some features are only available on certain device models. Warning icons and confirmation dialog boxes designate unsupported features.

Deployment considerations

- When you deploy an access control policy, its rules are not applied to existing connections. Traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic on an existing connection that could match a policy is omitted from the hit count. To have the policy rules effectively applied, clear the existing connection sessions, and then deploy the policy.
- For the system to affect traffic, you must deploy relevant configurations to managed devices using routed, switched, or transparent interfaces, or inline interface pairs.
- The system might prevent you from deploying inline configurations to passively deployed devices, including inline devices in tap mode. In other cases, the policy might deploy successfully, but attempts to block or alter traffic can produce unexpected results because passive devices do not enforce traffic actions. For example, the system might report multiple beginning-of-connection events for a connection that the policy attempts to block.

Best practices for creating, ordering, and implementing access control rules are detailed in Best practices for access control rules, on page 2058 and subtopics.

Best practices for ordering rules

Access control rules are matched top down, first match wins. Thus, to get the results you expect, you must pay careful attention to proper sequencing.

- In general, place top-priority rules that must apply to all traffic near the top of the policy. Always order rules to suit your organization's needs.
- Specific rules should come before general rules, especially when the specific rules are exceptions to general rules. Otherwise, traffic will match the general rule first and never hit the applicable specific rule.
- Rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number) should come as early as possible. Rules based on these criteria do not require inspection to identify matching connections.
- Whenever possible, put specific drop rules near the top of the policy. This ensures the earliest possible decision on undesirable traffic.
- URL filtering, application-based, and geolocation-based rules and others that require inspection should come after rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number), but before rules that specify file and intrusion policies.
- For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:
  - They include application conditions.
  - The traffic to be inspected is encrypted.
- Put URL filtering rules above application rules, and follow application rules with micro-application rules and Common Industrial Protocol (CIP) sub-classification application filtering rules.
- In general, rules with application conditions should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example. For more information, see Best practice for choosing between application matching and port matching, on page 2056 and Recommendations for application control, on page 2054.

- Rules that specify file policies and intrusion policies should come at the bottom of the rule order. These rules require resource-intensive deep inspection, and you should eliminate as many threats as possible using less-intensive methods first, for performance reasons, in order to minimize the number of potential threats that require deep inspection.

Application rule order

Place rules with application conditions lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

For more information, refer to Best practice for choosing between application matching and port matching, on page 2056 and Recommendations for application control, on page 2054.

Rule preemption

Rule preemption is a policy evaluation condition that occurs when a rule will never match traffic because a rule earlier in the evaluation order matches the traffic first. A rule's conditions govern whether it preempts other rules.

Rule preemption scenarios

Any type of rule condition can preempt a subsequent rule. This table shows common preemption scenarios:

Table 119: Rule preemption examples

Rule 1 (preempting)	Rule 2 (preempted)	Rule preemption scenario
Allow Admin users	Block Admin users	The second rule cannot block Admin traffic because the first rule allows it.
Allow Source Network VLAN 22-33	Block Source Network VLAN 27	The VLAN range in the first rule includes the VLAN in the second rule, so the first rule preempts the second. To block VLAN 27, you must move that rule above the one that allows VLAN 22-33.
Allow Source Network 10.4.0.0/16	Allow Source Network 10.4.0.0/16, VLAN 2	The VLAN range in the first rule includes the VLAN in the second rule, so the first rule preempts the second. To block VLAN 27, you must move that rule above the one that allows VLAN 22-33.
Allow Source Network 10.10.10.0/24 to URL www.netflix.com	Allow Source Network 10.10.10.0/24 to URL www.netflix.com	A rule also preempts an identical subsequent rule where all configured conditions are the same.

Best practice for rule actions and rule order

A rule's action determines how the system handles matching traffic. Improve performance by placing rules that do not perform or ensure further traffic handling before the resource-intensive rules that do. Then, the system can divert traffic that it might otherwise have inspected.

The following examples show how you might order rules in various policies, given a set of rules where none is more critical and preemption is not an issue.

If your rules include application conditions, also refer to Best practice for choosing between application matching and port matching, on page 2056.

Optimum order for access control rules

Intrusion, file, and malware inspection requires resources, especially if you use multiple custom intrusion policies and variable sets. Place access control rules that invoke deep inspection last.

1. Monitor—Rules that log matching connections, but take no other action on traffic. (However, see the important exception and caveat at Access control rule monitor action, on page 2046.)
2. Trust, Block, Block with reset—Rules that handle traffic without further inspection.
3. Allow, Interactive Block (no deep inspection)—Rules that do not inspect traffic further, but allow discovery.
4. Allow, Interactive Block (deep inspection)—Rules associated with file or intrusion policies that perform deep inspection for prohibited files, malware, and exploits.

Best practices for simplifying and focusing rules

Follow these best practices to simplify configuration and ensure optimal device performance and rule clarity:

Simplify rule conditions

Minimize individual rule criteria. Use as few elements in rule conditions as possible. For example, use IP address blocks rather than individual IP addresses when defining network conditions.

Avoid redundant conditions

If a single condition is enough to match the traffic you want to handle, do not add more. Redundant conditions can make configuration unnecessarily complex and may negatively impact device performance, especially during cluster or high-availability unit re-join.

For example:

- Use security zones that represent multiple interfaces carefully. If you specify source and destination networks as conditions, and these are enough to match the traffic you are targeting, then specifying a security zone is not required.
- If you want to match a set of internal interfaces to any destination on the Internet (for example), then simply use a source security zone that includes your internal interfaces. No network or destination interface criteria are needed.

Combine elements for organization, not performance

The use of network objects containing many individual IP addresses offers organizational benefits only. It does not improve device performance compared to listing those IP addresses individually.

For recommendations related to application detection, refer to Best practice for choosing between application matching and port matching, on page 2056.

Narrowly constrain resource-intensive rules, especially by interface

As much as possible, use rule conditions to narrowly define the traffic handled by resource-intensive rules. Focused rules are also important because rules with broad conditions can match many different types of traffic, and can preempt later, more specific rules. Examples of resource-intensive rules include:

- TLS/SSL rules that decrypt traffic—Not only the decryption, but further analysis of the decrypted traffic, requires resources. Narrow focus, and where possible, block or choose not to decrypt encrypted traffic.
- Access control rules that invoke deep inspection—Intrusion, file, and malware inspection requires resources, especially if you use multiple custom intrusion policies and variable sets. Make sure you only invoke deep inspection where required.
- If you specify security zones in a rule, the rule is deployed only to devices that have interfaces in the specified zone. So, if you want to focus a rule on only some devices that are assigned to the policy, ensure that you select a security zone that applies to the appropriate subset of devices. This ensures that unnecessary rules are not deployed to a device.

Maximum number of access control rules and intrusion policies

The maximum number of access control rules or intrusion policies that are supported by a device depends on many factors, including policy complexity, physical memory, and the number of processors on the device.

If you exceed the maximum supported by a device, you cannot deploy the access control policy and you must reevaluate.

Guidelines for intrusion policies

Consider these guidelines when configuring intrusion policies:

- In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.
- You may want to consolidate intrusion policies or variable sets so you can associate a single intrusion policy-variable set pair with multiple access control rules. On some devices you may find you can use only a single variable set for all your intrusion policies, or even a single intrusion policy-variable set pair for the whole device.

Access control rule management

Access control rule management is a security administration process that involves procedures for creating, modifying, and maintaining access control rules.

These topics explain how to manage access control rules.

Add an access control rule category

Categories help organize rules into logical groups so you can understand and manage your policy more easily.

You can divide an access control policy's **Mandatory** and **Default** rule sections into custom categories. After you create a category, you cannot move it, although you can delete it, rename it, and move rules into, out of, within, and around it. The Firewall Management Center assigns rule numbers across sections and categories.

Procedure

Step 1 In the access control policy editor, click **Add Category**.

If your policy already contains rules, you can click a blank area in the row for an existing rule to set the position of the new category before you add it. You can also right-click an existing rule and select **Insert new category**.

Step 2 Enter a **Name**.

Step 3 From the **Insert** drop-down list, choose where you want to add the category:

- To insert a category below all existing categories in a section, choose **into Mandatory** or **into Default**.
- To insert a category above an existing category, choose **above category**, then choose a category.
- To insert a category above or below an access control rule, choose **above rule** or **below rule**, then enter an existing rule number.

Step 4 Click **Apply**.

Step 5 Click **Save** to save the policy.

What to do next

You can now:

- Drag and drop rules into, or out of, the category.
- When creating a rule, select a category it should be part of.
- When editing a rule, reposition the rule to be in the category, or to move it out of the category.

Introduction to real-time Policy Analyzer and Optimizer

The real-time Policy Analyzer and Optimizer feature enhances access control rule management by providing immediate, real-time feedback on potential rule conflicts during configuration. Traditionally, users would add or modify access control rules without visibility into whether these new rules are shadowed or are made redundant by the existing rules in the access control policy.

While the on-demand Policy Analyzer and Optimizer within Security Cloud Control post-deployment analysis, the new real-time anomaly detection capability integrates proactive anomaly detection directly into the rule creation and editing workflow. This offers instant insights even as users create or update rules in Cloud-Delivered Firewall Management Center.

Benefits of real-time Policy Analyzer and Optimizer

The real-time Policy Analyzer and Optimizer feature addresses the common challenge of undetected rule anomalies, such as shadowing or redundancy, prior to deployment. These anomalies result in inefficient, cluttered rule sets having unnecessary or duplicate entries. By identifying conflicts even as rules are being added or modified, real-time anomaly detection helps administrators to:

- **Reduce rules set complexity**: Streamlines policies by highlighting and enabling the removal of redundant or shadowed rules.
- **Improve efficiency**: Eliminates the need for time-consuming post-deployment analysis, saving effort and accelerating policy deployment.
- **Maintain policy integrity**: Ensures that potential conflicts are identified and presented to users before they impact network security or performance.

How real-time Policy Analyzer and Optimizer works

Anomalies are detected and reported exclusively on the rule currently being created or edited, ensuring that existing rules are not independently flagged unless changes are made on them.

When you add or edit a rule, Cloud-Delivered Firewall Management Center checks for conflicts (shadowed or redundancy) based on the position of that rule:

- **Adding a rule after existing rules**: A new rule can be shadowed by rules that are already placed above it in the policy.
- **Adding a rule in the middle**: A new rule can be shadowed by the rules above it and may also make the rules positioned below it redundant.
- **Adding a rule at the beginning**: A new rule can make the rules positioned below it redundant.



Note When editing a rule, if changes introduce new anomalies, they will be reported, and the conflict type will be determined based on the edited rule's position.

Prerequisites for using real-time Policy Analyzer and Optimizer

To use the real-time Policy Analyzer and Optimizer, enable AgenticOps Insights in Security Cloud Control. For more information, see Enable AgenticOpsInsights.

Once AgenticOps Insights is enabled, the real-time Policy Analyzer and Optimizer will be enabled automatically. To view and manage the settings, navigate to **Insights & Reports** > **AgenticOps Insights** > **Settings** > **Policy Analyzer & Optimizer**.



Note You can choose to opt out of real-time Policy Analyzer and Optimizer if you no longer want to use this feature. Click the real-time Policy Analyzer and Optimizer toggle button and then click **Submit** to disable the feature.

Limitations of real-time Policy Analyzer and Optimizer

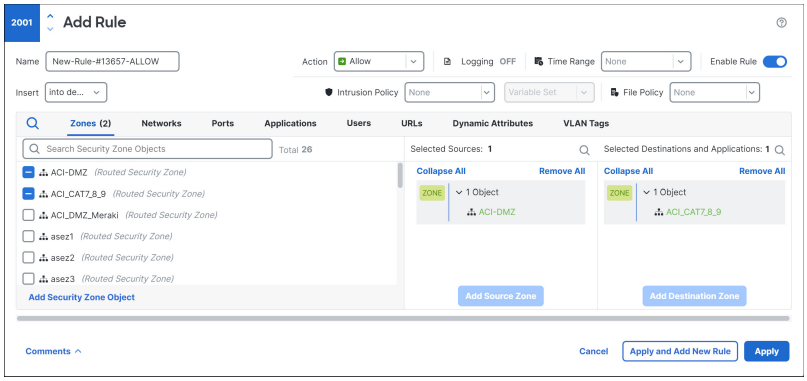
These are the limitations to consider when using the real-time Policy Analyzer and Optimizer feature:

- **Bulk operations**: Adding rules in bulk using APIs is not supported.
- **Rule type specificity**: Conflict detection is available only for access control rule types.
- **Workflow enabled**: Real-time Policy Analyzer and Optimizer is skipped if a workflow is enabled for rule management.
- **Stale policy cache**: Anomaly detection is skipped if another user has uncommitted changes. To proceed with real-time anomaly detection, users must first deploy or cancel pending changes, if any.
- **Device overrides**: Conflict detection considers only the primary value assigned to all object types when adding or modifying rules, and does not consider object overrides.
- **Disabled rules**: Disabled rules are excluded from conflict detection and do not appear as conflicts for other rules.

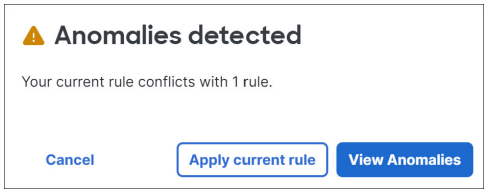
Manage rule anomalies

This section explains how the real-time Policy Analyzer and Optimizer feature in Cloud-Delivered Firewall Management Center automatically detects anomalies when you create a new rule or edit a rule. It also describes the options available in the interface for reviewing and resolving the detected anomalies. The same behavior applies when modifying an existing rule.

When you create a new rule, the following rule configuration dialog box is displayed:



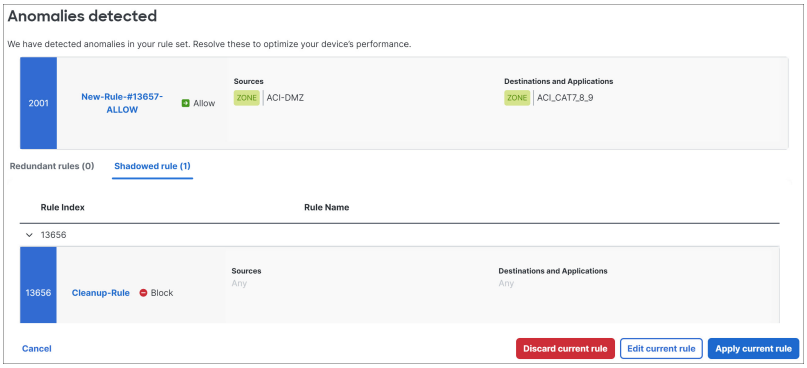
After defining the rule details and clicking **Apply**, the feature performs an automatic anomaly check on the new rule. If conflicts with existing rules are found, the **Anomalies detected** warning dialog box is displayed.



You can click one of the following options to:

- **Cancel**: Returns you to the rule editing interface to update the current rule.
- **Apply Current Rule**: Saves the rule despite the detected anomalies.
- **View Anomalies**: Opens a dialog box to display detailed information about the conflicts detected.

On clicking **View Anomalies**, the **Anomalies detected** dialog box is displayed, showing all the anomalies that have been detected in the context of to the current rule.



You can click one of the following options to:

- **Cancel**: Closes the anomalies dialog box.
- **Discard current rule**: Prompts a confirmation dialog box to discard the rule.
  - **Discard current rule**: The rule is discarded, and both the anomalies dialog box and the add rule dialog box are closed.
  - **Cancel**: The confirmation dialog box closes, returning you to the anomalies dialog box.
- **Edit current rule**: Closes the anomalies dialog box and returns you to the add rule dialog box to modify the rule details.
- **Apply current rule**: A dialog box is displayed prompting you to save the rule despite potential anomalies.
  - **Proceed**: The rule will be saved.
  - **Cancel**: The confirmation dialog box will close, returning you to the anomalies dialog box or the initial warning dialog box, depending on which is applicable.

Create and edit access control rules

Use access control rules to apply actions to specific traffic classes. Rules allow you to selectively allow desirable traffic and drop unwanted traffic.

Before you begin

If you have enabled the Policy Analyzer & Optimizer tool, your edits are evaluated as you make them. If anomalies are detected, you are notified during the edit and prompted to make changes when you click **Apply** to save the rule. You can view the anomalies, such as redundant and shadowed rules, and choose to edit or delete the rule. You can also proceed with saving the rule as is, to deal with the issue later.

Procedure

Step 1 In the access control policy editor, you have the following options:

- To add a new rule, click **Add Rule**.

- To edit an existing rule, click **Edit** ( ).
- To start with a copy of an existing rule, choose one of the following commands from the **More** ( ) menu.
  - **Copy Rule**, to copy the rule to the clipboard, so you can use the **Paste Above/Below** commands to place it anywhere in the same policy.
  - **Copy Rule to Different Policy**, to copy the rule to another access control policy. A dialog box opens so that you can select the policy and chose the rule placement.
  - **Clone Rule**, to create the copy immediately below the duplicated rule.
- To edit multiple rules, use the checkboxes to select multiple rules, then select **Edit** or another action from the **Select Action/Bulk Rule Actions** list next to the search box.
- To do inline editing, where you change the configuration of an object in a rule condition, right-click the value and choose **Edit**. You can also use the right-click menu to remove a item, add it to the filter, or copy the text or value.

If **View** ( ) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.

Step 2 If this is a new rule, enter a **Name**.

Step 3 Configure the rule components.

If you are bulk-editing multiple rules, only a subset of options are available.

Table 120: Access control rule components

Rule component	Configuration description
Position	Specify the rule position (**Insert** for new rules, the **Reposition** icon next to the rule name for existing rules); see Access control rule order, on page 2045.
Action	Choose a rule **Action**; see Access control rule actions, on page 2046.
Advanced Logging	Click **Advanced Logging** and choose the protocols for application-aware and protocol-aware event logs. See *About Application-Aware Event Logging* in the *Cisco Secure Firewall Management Center Administration Guide* for more information. **Note** The **Advanced Logging** option is available when the feature is enabled in the access control policy's advanced settings. See *Enable Application-Aware Event Logging* in the *Cisco Secure Firewall Management Center Administration Guide*.
Logging	Click **Logging** to specify options for connection logging and SNMP traps. .
Time Range	(Optional.) For Firewall Threat Defense devices, choose the days and times when the rule is applicable. If you do not choose an option, the rule is always active. For details, see Create a time range object, on page 3036.
Enable Rule	Whether the rule is active. Disabled rules are not applied to connections. You can disable rules to temporarily turn them off, for example, during troubleshooting.

Rule component	Configuration description
Security Protections (Deep inspection)	(Optional.) For Allow and Interactive Block rules, select options for **Intrusion Policy**, **Variable Set**, and **File Policy**. You can apply intrusion and file policies independently; you do not need to configure both.
Conditions	Select the objects you want to add to either source or destination, then click either **Add to Sources** or **Add to Destinations and Applications** to add matching conditions for connections. You can click a tab to restrict the list of available objects, for example, to Networks, Security Zones, Applications, and so forth. However, the sources and destination column always show all selected objects regardless of the tab you are on. See Access control rule conditions, on page 2068 for more information.
Comments	Open the comments list at the bottom of the dialog box, enter your comment, and click **Add Comment/Post** to add a comment.

Step 4 Click **Add** or **Apply** to save the rule.

Step 5 Click **Save** to save the policy.

What to do next

If you will deploy time-based rules, specify the time zone of the device to which the policy is assigned. See Configure a time zone for a device, on page 1304.

Deploy configuration changes.

Related Topics

Best practices for access control rules, on page 2058

Access control rule conditions

Access control rule conditions are traffic matching criteria that define the characteristics of the connections you want to target with each rule. Use the conditions precisely to fine-tune the rule to apply to all, and only, the traffic that should be handled by the rule. The following topics explain the match conditions that you can use.

Security/tunnel zone rule conditions

Security/tunnel zone rule conditions are traffic matching mechanisms that

- use security zones and tunnel zones to select traffic for a rule
- segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices
- allow you to identify tunneled traffic, such as GRE, that should be handled as a tunnel rather than apply access control rules to the encapsulated connections within the tunnel.

You can use security zones to control traffic by its source and destination interfaces. If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones for it to match the rule. Just as all interfaces in a security zone must be of the same type (all inline, passive, switched, or routed), all zones used

in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

When using tunnel zones, ensure that you have matching rules in the prefilter policy to associate tunneled traffic with the zone. Then, you can select the tunnel zone as a source zone in a rule; tunnel zones cannot be destinations. If you do not have prefilter rules to rezone the tunnels into the tunnel zone, an access control rule for the tunnel will never apply to any connections. You can specify destination security zones to target tunnels that leave the device through specific interfaces.

Security zone considerations

Consider the following when deciding on security zone criteria:

- Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.
- Access control rules generate ACL entries (ACEs) in the device configuration to provide early processing and drops whenever possible. If you specify security zones in rules, ACEs are created for each interface in the zone, which can greatly increase the size of the ACL. Excessively large ACLs generated from access control rules can impact system performance.

Network rule conditions

Network rule conditions are the network objects or geographical locations that define the network addresses or locations of the traffic.

Network criteria selection

- To match traffic from an IP address or geographical location, add the criteria to the Sources list.
- To match traffic to an IP address or geographical location, add the criteria to the Destinations list.
- If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from these tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control.

  Whenever possible, combine multiple network objects into a single object group. The system automatically creates an object group (during deployment) when you select more than one object (for source or destination separately). Selecting existing groups can avoid object group duplication and reduce the potential impact on CPU usage when there are a large number of duplicate objects.

  You can use objects that define the address using the fully-qualified domain name (FQDN); the address is determined through a DNS lookup. However, FQDN objects are not supported in these sections in access control policies: original client networks, SGT/ISE attributes, network analysis and intrusion policy, Security Intelligence, threat detection, elephant flow settings.

- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.

Note To ensure that you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

Traffic matching configuration:

- To match traffic from an IP address or geographical location, add the criteria to the Sources list.
- To match traffic to an IP address or geographical location, add the criteria to the Destinations list.
- If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

Original client in network conditions (filtering proxied traffic)

For some rules, you can handle proxied traffic based on the originating client. Use a source network condition to specify proxy servers, then add an original client constraint to specify original client IP addresses. The system uses a packet's X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header field to determine original client IP.

Traffic matches the rule if the proxy's IP address matches the rule's source network constraint, and the original client's IP address matches the rule's original client constraint. For example, to allow traffic from a specific original client address, but only if it uses a specific proxy, create three access control rules:

Access Control Rule 1: Blocks proxied traffic from a specific IP address (209.165.201.1)

Source Networks: 209.165.201.1
Original Client Networks: none/any
Action: Block

Access Control Rule 2: Allows proxied traffic from the same IP address, but only if the proxy server for that traffic is one you choose (209.165.200.225 or 209.165.200.238)

Source Networks: 209.165.200.225 and 209.165.200.238
Original Client Networks: 209.165.201.1
Action: Allow

Access Control Rule 3: Blocks proxied traffic from the same IP address if it uses any other proxy server.

Source Networks: any
Original Client Networks: 209.165.201.1
Action: Block

VLAN tag rule conditions

VLAN tags rule conditions are filtering mechanisms that

- control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic
- use the innermost VLAN tag to filter VLAN traffic, with the exception of prefilter policy which uses the outermost VLAN tag, and
- apply only to inline sets and do not match traffic on firewall interfaces.

VLAN tag rule conditions configuration details

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note this Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

User rule conditions

A user rule condition is a traffic matching mechanism that

- matches traffic based on whether the user initiating the connection or the group to which they belong
- can be configured for users in Microsoft Active Directory realms only, and
- enables policy enforcement based on user identity and group membership.

Special identities users

In addition to configuring users and groups for configured realms, you can set policies for the special identities users:

- Failed Authentication: User who failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.
- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified such as users who are not downloaded by a configured realm.

For access control rules only, you must first associate an identity policy with the access control policy as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

User rule condition example

Configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

Application rule conditions

Application rule conditions are application awareness and control mechanisms that

- identify and classify applications through IP traffic analysis
- enable control with organized filters
- require one enabled detector for each condition.

Application filters and control benefits

System-provided *application filters* help perform control by organizing applications according to basic characteristics, such as type, risk, business relevance, category, and tags. You can also create reuseable, user-defined filters using combinations of system-provided filters, or custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled, the system automatically enables all available system-provided detectors, or the most recently modified user-defined detector. For more information about application detectors, see Application detectors.

You can utilize both application filters and individually specified applications to ensure complete coverage. Understand the application's requirements to effectively order access control rules.

Application filters help you to quickly configure application control effectively. For example, you can easily use system-provided filters to create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

The system characterizes detected applications using several criteria which form effective filters.

Table 121: Application characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email applications are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.

Characteristic	Description	Example
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Configure application conditions and filters

To build an application condition or filter, choose the applications whose traffic you want to control from a list of available applications. It is recommended that you constrain the available applications using filters. You can use filters and individually specified applications in the same condition.

Before you begin

- Adaptive profiling must be enabled (its default state) in the access control policy's advanced settings for access control rules to perform application control.

Follow these steps to configure application conditions and filters:

Procedure

Step 1 In the access policy rule editor, click **Applications**.

Step 2 Find and choose the applications you want to add from the **Available Applications** list.

To constrain the applications displayed in **Available Applications**, choose one or more **Application Filters** or search for individual applications.

Tip

Click **Information** ( ) next to an application to display summary information and internet search links. A **lock** icon marks applications that the system can identify only in decrypted traffic.

When you choose filters, singly or in combination, the Available Applications list updates to display only the applications that meet your criteria. You can choose system-provided filters in combination, but not user-defined filters.

- Multiple filters for the same characteristic (risk, business relevance, and so on)—Application traffic must match only one of the filters. For example, if you choose both the medium and high-risk filters, the Available Applications list displays all medium and high-risk applications.
- Filters for different application characteristics—Application traffic must match both filter types. For example, if you choose both the high-risk and low business relevance filters, the Available Applications list displays only applications that meet both criteria.

Step 3 Click **Add Application** or **Add to Rule**, or drag and drop.

Tip
Before you add more filters and applications, click **Clear Filters/Selection** to clear your current choices.

Step 4 If destination ports are not specified on the Ports tab, select the **Port** for the applications.

The **Port** specification applies to all the applications in the list; it does not apply to any filters. You cannot specify different port options for each application. If you already specified destination ports on the Ports tab, that selection is reflected in the application list, and you cannot make the selection.

If you have not already specified destination ports for the rule, choose one of the following:

- **Application Default**—The system looks at the default port only when evaluating traffic for an application match. Click the info button (i) to see what is considered the default for each application.

  When editing a rule, if the rule already has source ports, and you add an application that has an incompatible port and select Application Default port, you are warned. If you elect to use the application default port, the incompatible source port is removed.

- **Any**—The system does not restrict the application identification based on the port used in the connection. If there are source ports specified on the Port tab, you might need to select this option to be able to save the rule.

Regardless of the option you select, if you subsequently configure any destination value on the Ports tab, the ports specified on the Ports tab override your selection here, and change the application port to Any. If you want to limit the rule to one, multiple, or a range of ports that are not the application default port, use the Ports tab.

This option does not apply to custom detectors.

Note
If the access control policy is assigned to devices running 7.x versions, the rule is handled differently if you have included applications that have different default ports. When you select to apply the rule to the default ports only, then for 7.7 systems, the rule is split into separate rules, one per application. For systems running versions 7.3-7.6, the rule is kept as a single rule, but the ports page is updated with each of the default ports. If you are managing devices running down-level software versions, and you want to use default port matching, we recommend that you limit the application's list in each rule to applications that have the same default port.

Step 5 Save or continue editing the rule or configuration.

What to do next

Deploy configuration changes.

Port, protocol, and ICMP code rule conditions (concept)

A port, protocol, and ICMP code rule condition is a traffic matching criterion that

- matches traffic based on source and destination ports
- supports TCP, UDP, ICMP, and other protocols that do not use ports, and
- controls traffic flow in access control and other policy rules.

Port condition types

Port conditions match traffic based on the source and destination ports. Depending on the rule type, "port" can represent any of these types:

- TCP and UDP—Control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP (6) /22.
- ICMP—Control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—Control traffic using other protocols that do not use ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best practices for port-based rules

Specifying ports is the traditional method used to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Use application filtering criteria when possible; however, note that application filtering is unavailable in prefilter rules.

Use application filtering for FTP and similar applications that open separate channels for control and data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly and could result in blocking desirable connections.

Using source and destination port constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching non-TCP traffic with port conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although configuring port conditions can match non-TCP traffic, restrictions apply:

- Access control rules—For Classic devices, match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, add only network-based conditions: zone, IP address, port, and VLAN tag. The system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support only TCP port conditions.
- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies, sent in response to ICMP echo requests, are ignored. To match any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Recommendation: use URL conditions for website access control

Use URL conditions to control the websites that users on your network can access.

For complete information, see URL filtering with category and reputation, on page 2086.

Dynamic attributes rule conditions

Dynamic attributes rule conditions allow you to specify access control rules based on real-time properties such as network objects, device types, locations, and security tags. Dynamic attributes can be used as source criteria and destination criteria in access control rules. These criteria help match traffic more flexibly and accurately.

Dynamic attributes configuration and behavior

You can use these types of dynamic attributes to match connections in DNS rules:

- Dynamic objects (Source or destination): The dynamic attributes connector collects data such as networks and IP addresses from cloud providers for use in access control rules. For more information, refer to About the Dynamic Attributes Connector , on page 2621 .
- SGT objects (Source only): These contain tags defined manually or in Cisco ISE. For more information, refer to Source and destination security group tag (SGT) matching, on page 2374 and Security group tags, on page 2972 .
- Location IP objects (Source only): Defined by Cisco ISE.
- Device type objects (Source only): Defined by Cisco ISE, also referred to as endpoint profile objects.

When you configure dynamic attributes for a DNS rule, objects of the same type within a source or destination list use OR logic, while objects of different types use AND logic. For example, if you select both a security group tag, and a dynamic object that lists IP addresses, the rule matches if traffic with the tag originates from (or is destined to) one of those IP addresses.

When you open the **Dynamic Attributes** tab, all **Dynamic Objects** and **Security Group Tags** appear in the list. To remove objects from the list, deselect the options. To find an object, type in the search box.

Time and Day Rule Conditions

You can specify a continuous time range or a recurring time period.

For example, a rule can apply only during weekday working hours, or every weekend, or during a holiday shutdown period.

Time-based rules are applied based on the local time of the device that processes the traffic.

Time-based rules are supported only on Firewall Threat Defense devices. If you assign a policy with a time-based rule to a different type of device, the time restriction associated with the rule is ignored on that device. You will see warnings in this case.

Enable and disable access control rules

Enable or disable access control rules to control whether they are used to evaluate network traffic.

When you create an access control rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing rules in an access control policy, disabled rules are grayed out but can still be modified. You can also change a rule's status in the rule editor.

Procedure

Step 1 In the access control policy editor, right-click the rule you want to change and choose a rule state.

If **View** ( ) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.

Step 2 Click **Save**.

What to do next

- Deploy configuration changes.

Copy access control rules from one access control policy to another

You can copy access control rules from one access control policy to another. You can copy the rules either to the **Default** section or the **Mandatory** section of the access control policy.

All the settings of the copied rules, except the comments, are retained in the pasted version

Procedure

Step 1 Do one of the following:

- To copy a single rule, right-click the rule and select **Copy Rule to Different Policy**.
- To copy multiple rules, select their checkboxes, then select **Copy Rules to Different Policy** from the **Select Bulk Action** menu.

Step 2 Select the destination access control policy from the **Access Policy** drop-down list.

Step 3 From the **Place Rules** drop-down list, choose where you want to position the copied rules. You can place them at the top or bottom of either the Mandatory or Default sections.

Step 4 Click **Copy**.

What to do next

- Deploy configuration changes.

Move access control rules to a prefilter policy

You can move access control rules from an access control policy to the associated non-default prefilter policy.

You must first apply a user-defined prefilter policy to the access control policy. The access control rules cannot be moved to the default prefilter policy because the default prefilter policy cannot have rules.

Before you begin

Note these conditions before you proceed:

- When moving an access control rule to a prefilter policy the layer 7 (L7) parameters, such as application or URL filterin, in the access control rule cannot be moved. The L7 parameters are dropped during the operation.
- The comments in the access control rule configuration are lost after moving the rule. However, a new comment is added in the moved rule mentioning the source access control policy.
- You cannot move access control rules with **Monitor** set as the **Action** parameter.
- The **Action** parameter in the access control rule is changed to a suitable action in the prefilter rule when moved. To know what each action in the access control rule maps to, see this table:

Action in the access control rule	Action in the prefilter rule
Allow	Analyze
Block	Block
Block with reset	Block
Interactive Block	Block
Interactive Block with reset	Block
Trust	Fastpath

- Similarly, based on the action configured in the access control rule, the logging configuration is set to an appropriate setting after the rule is moved, as mentioned in this table.

Action in the access control rule	Enabled logging configurations in the prefilter rule
Allow	None of the check boxes are checked.
Block Block with reset Interactive Block Interactive Block with reset	• Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Trust	• Log at Beginning of Connection • Log at End of Connection • Event Viewer • Syslog Server • SNMP Trap

- While moving rules from the source policy, if another user modifies those rules, you will see get a message. You may continue with the process after refreshing the page.

Procedure

Step 1 Do one of the options:

- To move a single rule, right-click the rule and select **Move Rule to Prefilter Policy**.
- To move multiple rules, select their checkboxes, then select **Move Rules to Prefilter Policy** from the **Select Bulk Action** menu.

Step 2 From the **Place Rules** drop-down list, choose where you want to position the moved rules, either **At the bottom** or **At the top**.

Step 3 Click **Move**.

What to do next

- Deploy configuration changes.

Position an access control rule

Position rules correctly within an access control policy to ensure the first match wins behavior produces the intended security outcomes.

You can move an existing rule within an access control policy, or insert new rules in a desired location. When you add or move a rule to a category, the system places it last in the category.

Before you begin

Review rule order guidelines in Best practices for access control rules, on page 2058.

Procedure

Step 1 Do one of these:

- New rule—Insert a new rule by hovering over the line between the existing rules, and clicking **Add Rule**. The location is selected in the **Insert** box in the Add Rule dialog box; you can select a different rule to adjust the location. You can also select **Add Rule Above** or **Add Rule Below** from the right-click menu.
- Existing rules when viewing the rule table—Click and drag the rule to the new position. This action is final, you are not prompted to confirm it.
- Existing rules when viewing the rule table—Right-click a single rule and select **Reposition Rule**. To move multiple rules as a group, select their checkboxes, then select **Reposition Rules** from the **Select Bulk Action** menu. You are prompted for where you want the rule to go.
- Existing rule when editing the rule—Click the **Reposition Rule** icon next to the rule name.

Step 2 When editing or repositioning the rule, choose where you want to move or insert the rule, then click **Move**, **Confirm**, or **Reposition** (depending on what you are doing).

- Choose **into Mandatory** or **into Default**.
- Choose a **into Category**, then choose the category.
- Choose **above rule** or **below rule**, then select the rule.

Step 3 Save the rule if you are editing it.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Add comments to an access control rule

Adding comments to access control rules helps you keep track of the purpose of a rule, and changes made to the rule.

When you create or edit an access control rule, you can add a comment. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change. You can display a list of all comments for a rule along with the user who added each comment and the date the comment was added.

When you save a rule, all comments made since the last save become read-only.

To search access control rule comments, use the **Search Rules** bar on the rule listing page.

Procedure

Step 1 In the access control rule editor, click **Comments**.

Step 2 Enter your comment and click **Add Comment**. You can edit or delete this comment until you save the rule.

Step 3 Save the rule.

Examples for access control rules

These topics provide examples of access control rules.

How to control access using security zones

Configure security zones and access control rules to inspect incoming traffic for intrusions and malware while allowing unrestricted internet access for internal hosts.

Consider a deployment where you want hosts to have unrestricted access to the internet, but you nevertheless want to protect them by inspecting incoming traffic for intrusions and malware.

First, create two security zones: Internal and External. Then, assign interface pairs on one or more devices to those zones, with one interface in each pair in the Internal zone and one in the External zone. Hosts connected to the network on the Internal side represent your protected assets.

Note You are not required to group all internal (or external) interfaces into a single zone. Choose the grouping that makes sense for your deployment and security policies.

Then, configure an access control rule with a destination zone condition set to Internal. This simple rule matches traffic that leaves the device from any interface in the Internal zone. To inspect matching traffic for intrusions and malware, choose a rule action of **Allow** , then associate the rule with an intrusion and a file policy.

Control application usage

Block web-based applications, particularly anonymizers or proxies, by creating access control rules that identify and restrict traffic based on application attributes

Enterprise networks increasingly rely on web-based delivery platforms and rich media applications, often using the same web protocols. Firewall Threat Defense inspects connections to determine the application, enabling access control rules targeting applications rather than just targeting specific TCP/UDP ports.

Although you can select specific applications to allow or block, you can also write rules based on type, category, tag, risk, or business relevance. For example, you could create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is blocked.

Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

In this use case, we will block any application that belongs to the **anonymizer/proxy** category.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control** and edit the access control policy.

Step 2 Click **Add Rule** and configure the rule for application control.

a) Name the rule (for example, Block_Anonymizers).

b) For **Action**, select **Block**.

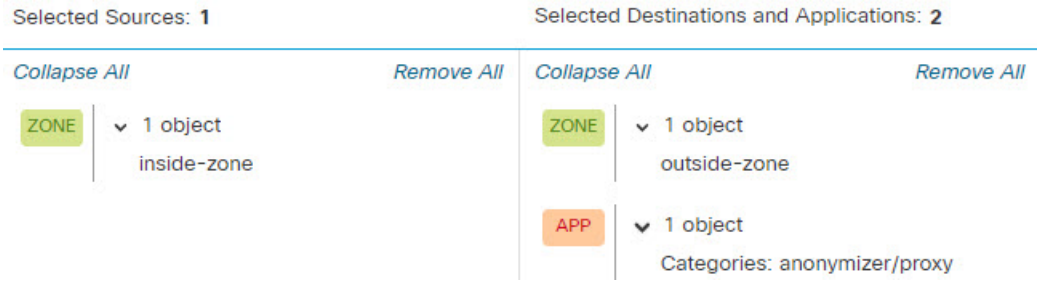


c) If you have configured security zones and want to restrict traffic from inside to outside, click the **Zones** tab, and choose your inside zone as the source zone, and the outside zone as the destination zone.

d) Click the **Applications** tab, select the applications to match, and click **Add Application**.

As you select criteria, such as category and risk level, the list to the right of the criteria updates to show exactly which applications match the criteria. The rule you are writing applies to these applications.

For purposes of this example, choose **anonymizers/proxy** from the **Categories** list and add it to **Destinations and Applications**. The match criteria should now look similar to the following graphic.



Note
For **Port**, leave the default, Application Default. This ensures that the system looks only at the default ports used for a selected application, rather than all ports.

e) Click **Logging** next to the rule action, and enable logging at the start of the connection. You can select a syslog server if you use one.

You must enable logging to get information about any connections blocked by this rule.

Step 3 Move the rule so that it comes after any rules that use protocol and port criteria only, but that would not allow traffic that should be blocked by the application rule.

Matching applications requires Snort inspection. Because Snort inspection is not needed by rules that use protocol and port only, you can improve system performance by grouping these simple rules at the top of the access control policy as much as possible.

Step 4 Deploy the changes.

You can use the application rule hit counts and analysis dashboards to see how this rule is performing and how often users try these applications.

Block threats

You can implement next generation Intrusion Prevention System (IPS) filtering by adding intrusion policies to your access control rules. Intrusion policies analyze network traffic, comparing the traffic contents against known threats. If a connection matches a threat you are monitoring, the system drops the connection, thus preventing the attack.

All other traffic handling occurs before network traffic is examined for intrusions. By associating an intrusion policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy.

You can configure intrusion policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic. In addition, you can configure an intrusion policy as the default action if you do not want to use a simple block.

Security Intelligence policies allow you to preemptively block traffic to or from known bad IP addresses and URLs. This task demonstrates how to add an intrusion policy that allows traffic from the internal 192.168.1.0/24 network to external destinations, assuming block rules exist to eliminate unwanted connections, and adds Security Intelligence to block known threats.

Before you begin

- Apply the IPS license to any managed device that uses this rule.
- Ensure you have created security zones for inside and outside interfaces.
- Create the network object for the inside network.

Procedure

Step 1 Create the access control rule that applies the intrusion policy.

a) Edit the access control policy and add a new rule.

b) Name the rule (for example, Inside_Outside) and set the rule action to **Allow**.



c) Choose **Balanced Security and Connectivity** as the **Intrusion Policy**. You can either accept the default variable set or select your own if you want to customize it.

The **Balanced Security and Connectivity** policy is appropriate for most networks. It provides a good intrusion defense without being overly aggressive, which has the potential of dropping traffic that you might not want to be dropped. If you determine that too much traffic is getting dropped, you can ease up on intrusion inspection by selecting the **Connectivity over Security** policy.

If you need to be aggressive about security, try the **Security over Connectivity** policy. The **Maximum Detection** policy offers even more emphasis on network infrastructure security with the potential for even greater operational impact.

If you create your own custom policy, you can select that one instead.

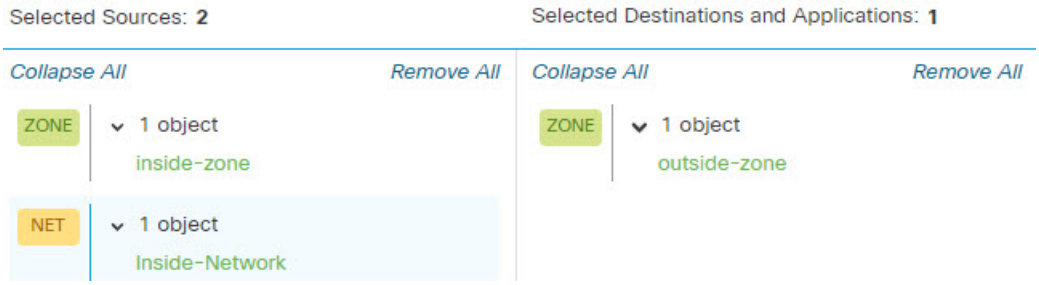
A discussion of variable sets is beyond the scope of this example. Please read the chapters on intrusion policy for detailed information about variable sets and custom policies.



d) Select the **Zones** tab, and add your inside security zone to the source criteria, and outside zone to the destination criteria.

e) Select the **Networks** tab, and add the network object that defines your inside network to the source criteria.

The match criteria should look similar to the following:



f) Click **Logging** and enable logging at the beginning or end of the connection, or both, as desired.

g) Click **Apply** to save the rule, and then **Save** to save the updated policy.
h) Move the rule to the appropriate location in the access control policy.

Step 2 Configure the Security Intelligence policy to preemptively drop connections with known bad hosts and sites.

By using Security Intelligence to block connections with hosts or sites that are known to be threats, you save your system the time needed to do deep packet inspection to identify threats in each connection. Security Intelligence provides an early block of undesirable traffic, leaving more system time to handle the traffic you really care about.

a) While editing the access control policy, click the **Security Intelligence** link in the packet path.

The link includes two policies: the DNS policy at the top, and the Security Intelligence (Network and URL) at the bottom. In this example, we are configuring the Network and URL lists. By default, these lists already include the global block and do not block lists, but these lists are empty by default until you add items to them.

b) With **Networks** selected and the **Any** security zone selected, scroll down in the Networks list until you get to the global feeds, and the first Security Intelligence category (probably Attackers). Click Attackers, then scroll to the end of the categories (probably Tor_exit_node), and Shift+Click to select all of the categories. Click **Add To Block List**.
c) Select the **URL** tab, and **Any** security zone, and use Shift+Click to select URL versions of the same categories. Click **Add To Block List**.
d) Click **Save** to save the policy.
e) As necessary, you can add network and URL objects to the block or do not block lists.

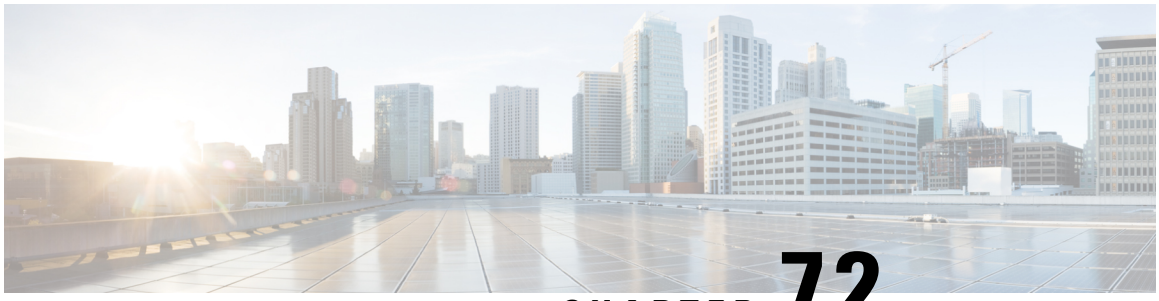
The **Do Not Block** lists are not real "allow" lists. They are exception lists. If an address or URL in the exception list also appears in the blocked list, the connection for the address or URL is allowed to pass on to the access control policy.

This way, you can block a feed, but if you later find that a desirable address or site is being blocked, you can use the exception list to override that block without needing to remove the feed entirely.

Keep in mind that these connections are subsequently evaluated by access control, and if configured, an intrusion policy. Thus, if any connections do contain threats, they can be identified and blocked during intrusion inspection.

Use the events and dashboards to determine what traffic is actually being dropped by the policy, and whether you need to add addresses or URLs to the **Do Not Block** lists.

Step 3 Deploy your changes.



CHAPTER 72

URL Filtering Rules

You can implement URL filtering using access control rules.

- URL filtering, on page 2085
- Best practices for URL filtering, on page 2088
- License requirements for URL filtering, on page 2093
- Requirements and prerequisites for URL filtering, on page 2094
- Configure URL filtering with category and reputation, on page 2094
- Manual URL filtering, on page 2101
- HTTP response pages, on page 2103
- Configure URL filtering health monitors, on page 2107
- Dispute URL category and reputation, on page 2107
- Take action if the URL category set changes, on page 2108
- Troubleshoot URL filtering, on page 2110

URL filtering

URL filtering is a security feature that

- controls the websites that users on your network can access,
- provides category and reputation-based filtering with a URL Filtering license, and
- enables manual specification of individual URLs, groups of URLs, and URL lists with any license.

URL filtering types

Use URL filtering to control website access through these methods:

- Category and reputation-based URL filtering—With a URL Filtering license, you can control access to websites based on the URL's general classification (category) and risk level (reputation). This is the recommended option.
- Manual URL filtering—With any license, you can manually specify individual URLs, groups of URLs, and URL lists and feeds to achieve granular, custom control over web traffic. For more information, see Manual URL filtering, on page 2101.

See also Security Intelligence, on page 2027, a similar but different feature for blocking malicious URLs, domains, and IP addresses.

The URL filtering feature is applicable only when you use a web browser to browse websites via the HTTP/HTTPS protocol. For other applications, you must use an FQDN-based access control rule instead of URL filtering. See Access Control Rules for more information.

URL filtering with category and reputation

URL filtering with category and reputation is a web security mechanism that

- controls access to websites based on the category and reputation of requested URLs,
- uses category classifications to group URLs by type (for example, ebay.com belongs to the Auctions category, and monster.com belongs to the Job Search category),
- assesses reputation levels from Unknown risk (level 0) or Untrusted (level 1) to Trusted (level 5) to determine how likely the URL is to be used for purposes that might be against the security policy for your organization, and
- allows URLs to belong to multiple categories simultaneously.

Reputation is determined by the hostname, e.g. all pages on www.cisco.com are considered to have the same reputation.

Benefits of category and reputation-based URL filtering

URL categories and reputations help you quickly configure URL filtering. For example, you can use access control to block untrusted URLs in the Hacking category. You can also use QoS to limit the rate of traffic from sites in the Streaming Video category. There are also categories for types of threats, such as a Spyware and Adware category.

Using category and reputation data simplifies policy creation and administration. It grants you assurance that the system controls web traffic as expected. Because Cisco continually updates its threat intelligence with new URLs, as well as new categories and risk levels for existing URLs, the system uses up-to-date information to filter requested URLs. Sites that (for example) represent security threats, or serve undesirable content may appear and disappear faster than you can update and deploy new policies.

These examples show how the system can adapt:

- If an access control rule blocks all gaming sites, as new domains are registered and classified as Games, the system can block those sites automatically. Similarly, if a QoS rule rate-limits traffic from all streaming video sites, the system can automatically limit traffic to new streaming video sites.
- If an access control rule blocks all malware sites and a shopping page gets infected with malware, the system can recategorize the URL from Shopping to Malware Sites and block that site.
- If an access control rule blocks untrusted social networking sites and somebody posts a link on their profile page that contains links to malicious payloads, the system can change the reputation of that page from Favorable to Untrusted and block it.

Limitations exist for category-based filtering in decryption policy Do Not Decrypt rules.

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from

its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security.

Note Don't confuse URL filtering with application detection, which relies on reading some of the packet from a website to determine more specifically what it is (for example, Facebook Message or Salesforce). For more information, see Recommendations for application control, on page 2054.

For more information, refer to Use categories in URL filtering, on page 2092.

URL category and reputation descriptions

This reference lists URL category and reputation level descriptions for security management.

Category descriptions

You can find a description of each URL category at https://www.talosintelligence.com/categories.

Click **Threat Categories** to view those categories.

Reputation level descriptions

Refer to Web reputation levels for information.

URL filtering data from the Cisco cloud

URL filtering data from the Cisco cloud is a threat intelligence dataset that

- allows traffic handling based on a website's general classification, or *category*, and risk level, or *reputation*,
- is automatically enabled when a URL Filtering license is added, and
- provides real-time threat intelligence evaluation for websites not found in local cache or datasets.

URL filtering data evaluation process

By default, when users browse to a URL whose category and reputation are not in a local cache of previously accessed websites, the system submits it to the cloud for threat intelligence evaluation and adds the result to the cache.

Optionally, you can use a local URL dataset of categories and reputations, which can make web browsing faster. When you enable or re-enable URL filtering, the Cloud-Delivered Firewall Management Center automatically queries Cisco for URL data and pushes the dataset to managed devices. Then, when users browse to an URL, the system checks the local dataset and the cache for category and reputation information before submitting it to the cloud for threat intelligence evaluation. To see your options for using the local dataset, including how to disable individual cloud lookups altogether, see URL filtering options, on page 2096.

Automatic updates of URL data is enabled by default. We strongly recommend that you do not disable these updates.

Note The Secure Firewall Threat Defense 200 series device does not store URL category and reputation data locally. It relies entirely on browser cache and cloud-based lookups for URL filtering and reputation evaluation.

The set of URL categories may change periodically. When you receive a change notification, review your URL filtering configurations to make sure traffic is handled as expected. For more information, see Take action if the URL category set changes, on page 2108.

Best practices for URL filtering

Filter by category and reputation

Follow the instructions in Configure URL filtering with category and reputation, on page 2094.

Configure your policy to inspect packets that must pass before a URL can be identified

The system cannot filter URLs before:

- A monitored connection is established between a client and server.
- The system identifies the DNS, HTTP or HTTPS application in the session.
- The system identifies the requested domain or URL (for encrypted sessions, from a non-encrypted domain name, the ClientHello message or the server certificate).

This identification should occur within three to five packets, or after the server certificate exchange in the TLS/SSL handshake if the traffic is encrypted.

Important! To ensure that your system examines these initial packets that would otherwise pass, see Inspection of packets that pass before traffic is identified, on page 2918 and subtopics.

If early traffic matches all other rule conditions but identification is incomplete, the system allows the packet to pass and the connection to be established (or the TLS/SSL handshake to complete). After the system completes its identification, the system applies the appropriate rule action to the remaining session traffic.

Block threat categories

Be sure that your policies specifically address Threat categories, which identify known malicious sites. Do this in addition to blocking sites with poor reputations.

For example, to protect your network from malicious sites, you must block all Threat categories. Additionally, Talos recommends that you block only sites with Poor category. You can block questionable reputations if you have an aggressive security posture, but this may result in a higher amount of false positives.

For specifics, see **Threat Categories** at the URL in URL category and reputation descriptions, on page 2087.

Position URL rules properly and consider rule order

Position URL rules after all other rules that *must* be hit.

URLs can belong to more than one category. It is possible to want to allow one category of websites and block another—whether explicitly or by relying on the default action. In this case, make sure you create and order URL rules so you get the desired effect, depending on whether the allow or the block should take precedence.

For additional guidelines for rules, see the following topics: Best practices for access control rules, on page 2058.

Handle uncategorized or reputationless URLs appropriately

When you build a URL rule, you first choose the category you want to match. If you explicitly choose **Uncategorized** URLs, you cannot further constrain by reputation.

Uncategorized URLs with Untrusted reputation are handled by the **Malicious Sites** category. If you want to block uncategorized sites with any other reputation level (such as Questionable), you must block all uncategorized sites.

After selecting a category and a reputation level, you can optionally select **Apply to unknown reputation**. For example, you can create a rule that applies to sites with Untrusted, Questionable, and unknown reputations.

You cannot manually assign categories and reputations to URLs, but in access control and QoS policies, you can manually block specific URLs. See Manual URL filtering, on page 2101. See also Dispute URL category and reputation, on page 2107.

URL filtering for encrypted web traffic

When performing URL filtering on encrypted web traffic, the system:

- (If DNS filtering is enabled) Checks to see if the system has previously seen the originating domain or the domain is in the local reputation database, and if so, takes action based on the reputation and category of the domain. Otherwise, the system processes the traffic based on your configurations for encrypted traffic, even if **Retry URL cache miss lookup** is enabled in the access control policy's advanced settings.
- Disregards the encryption protocol; a rule matches both HTTPS and HTTP traffic if the rule has a URL condition but not an application condition that specifies the protocol.
- Does not use URL lists. You must use URL objects and groups instead.
- Matches HTTPS traffic based on the subject common name in the public key certificate used to encrypt the traffic, and also evaluates the reputation of any other URLs presented at any time during the transaction, including the post-decryption HTTP URL.
- Disregards subdomains within the subject common name.
- Does not display an HTTP response page for encrypted connections blocked by access control rules (or any other configuration); see Limitations to HTTP response pages, on page 2103.

URL filtering and TLS server identity discovery considerations

The latest version of the Transport Layer Security (TLS) protocol 1.3, defined by RFC 8446, is the preferred protocol for many web servers to provide secure communications. Because the TLS 1.3 protocol encrypts the server's certificate for additional security, and the certificate is needed to match application and URL filtering criteria in access control rules, the Firepower System provides a way to extract the server certificate *without* decrypting the entire packet.

Access control policy advanced settings offer an **Early application detection and URL categorization** option for TLS Server Identity Discovery.

We strongly recommend enabling it for any traffic you want to match on application or URL criteria, especially if you want to perform deep inspection of that traffic. An decryption policy is not required because *traffic is not decrypted* in the process of extracting the server certificate.

Note

- TLS server identity discovery *cannot* be used with any of the following:
  - STARTTLS traffic
  - The HTTP CONNECT method
  - Traffic that is already being decrypted by another device on the network
- Because the certificate is decrypted, TLS server identity discovery can reduce performance depending on the hardware platform.
- TLS server identity discovery is not supported in inline tap mode or passive mode deployments.
- Enabling TLS server identity discovery is supported on any Secure Firewall Threat Defense Virtual version deployed in AWS. TLS server identity is incompatible with certain deployment scenarios that involve Gateway Load Balancer (GWLB). Whether the device initiates a server certificate probe depends on its deployment method with GWLB and the traffic direction.
  - Geneve single-arm flow: The probe *is not initiated* for either east–west or north–south traffic. A connection event **PROBE_FLOW_DROP_BYPASS_PROXY** increments every time the device attempts to extract the server certificate.
  - Geneve dual-arm flow: The probe *is initiated only* for north–south traffic.
- TLS Server Identity Discovery also operates on TLS 1.2 sessions.

For more information, see TLS server identity discovery, on page 2692.

HTTP/2 URL extraction limitations

The system can extract HTTP/2 URLs from TLS certificates, but not from a payload.

Use manual URL filtering effectively

- Specify URLs using a custom Security Intelligence list or feed object. Do not use a URL object or directly enter a URL into the rule. For details, see Manual URL filtering options, on page 2102.
- If you manually filter specific URLs using URL objects or by entering URLs directly into the rule, carefully consider other traffic that might be affected. To determine whether network traffic matches a URL condition, the system performs a simple substring match. If the requested URL matches any part of the string, the URLs are considered to match.
- If you use manual URL filtering to create exceptions to other rules, position the specific rule with the exceptions above the general rule that would otherwise apply.

Search query parameters in URLs are not used for matching

The system does not use search query parameters in the URL to match URL conditions. For example, consider a scenario where you block all shopping traffic. In that case, using a web search to search for amazon.com is not blocked, but browsing to amazon.com is.

Maintain default access control policy advanced settings

Unless you have good reason to change them, leave these settings on the Advanced tab of your access control policy at their defaults:

- **General Settings**:
  - **Retry URL cache miss lookup**
  - **Enable reputation enforcement on DNS traffic**

For descriptions of these settings, refer to General settings, on page 2691.

Reduce unnecessary usage of connection table space and NAT table slots

Consider a scenario in which URL filtering is enabled and any traffic that does not conform with the URL filtering rules is blocked. In such a scenario, the threat defence device analyses the initial packets in the traffic flow to extract URL information. After analysis, if the flow does not match any URL filtering rules, the flow is blocked. This means that the packets are dropped by the Threat Defense device. However, the associated connection remains in the connection table. This behaviour results in unnecessary use of connection table space and, if NAT is configured, NAT table slots, which can lead to increased consumption of network resources.

There are different ways to prevent this type of scenario. Since most of the traffic will be TCP, you can use the 'block with reset' option for TCP traffic instead of using 'block'. This approach reduces the usage of connection table space and NAT table slots.

URL filtering in high availability deployments

For guidelines for URL filtering with Firepower Management Centers in high availability, see *URL Filtering and Security Intelligence* in the *Cisco Secure Firewall Management Center Administration Guide*.

Memory limitations for selected device models

Device models with less memory store less URL data locally, and the system may therefore check the cloud more frequently to determine category and reputation for sites that are not in the local database.

Lower-memory devices include

- Firepower 1010
- Firewall Threat Defense Virtual with 8 GB of RAM

HTTPS traffic filtering

HTTPS traffic filtering is a security mechanism that

- determines the requested URL based on information passed during the TLS/SSL handshake using the subject common name in the public key certificate,
- disregards subdomains within the subject common name, requiring manual filtering to use top-level domains only, and
- does not support URL lists, requiring the use of URL objects and groups instead.

HTTPS filtering characteristics

There are several key differences between HTTPS filtering and HTTP filtering:

- Do not include subdomain information when manually filtering HTTPS URLs in access control or QoS policies. For example, use example.com instead of www.example.com.
-

Tip In an decryption policies, you can handle and decrypt traffic to specific URLs by defining a distinguished name decryption policy rule condition. The common name attribute in the subject distinguished name of a certificate contains the site's URL. Decrypting HTTPS traffic allows access control rules to evaluate the decrypted session, which improves URL filtering.

Traffic control by encryption protocol:

The system disregards the encryption protocol (HTTP vs HTTPS) when performing URL filtering in access control or QoS policies. This occurs for both manual and reputation-based URL conditions. In other words, URL filtering treats traffic to these websites identically:

- http://example.com/
- https://example.com/

To configure a rule that matches only HTTP or HTTPS traffic, add an application condition to the rule. For example, you could allow HTTPS access to a site while disallowing HTTP access by constructing two access control rules, each with an application and URL condition.

Application-specific URL filtering rules

The first rule allows HTTPS traffic to the website:

- Action: Allow
- Application: HTTPS
- URL: example.com

The second rule blocks HTTP access to the same website:

- Action: Block
- Application: HTTP
- URL: example.com

Use categories in URL filtering

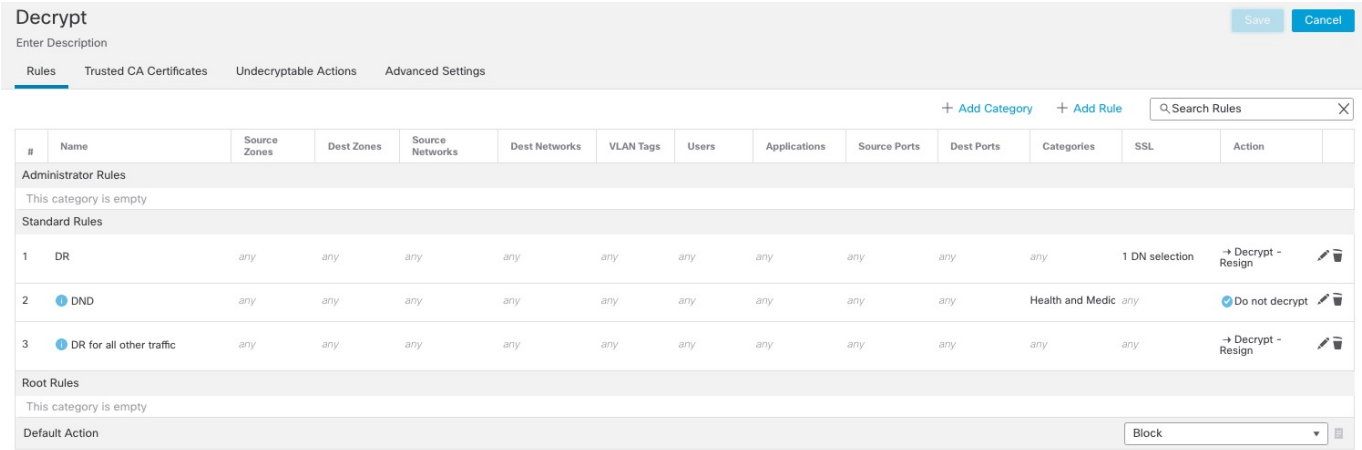
Limitations of categories in Do Not Decrypt rules

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security. While we strive to continuously update and improve URL filtering categories, it is not an exact

science. Some websites are not categorized at all and it's possible some websites might be improperly categorized.

Avoid overusing categories in do not decrypt rules to avoid decrypting traffic without a reason; for example, the Health and Medicine category includes the WebMD website, which does not threaten patient privacy.

Following is a sample decryption policy that can prevent decryption for websites in the Health and Medicine categories but allows decryption for WebMD and everything else. General information about decryption rules can be found in Guidelines for using TLS/SSL decryption, on page 2188.



Note Don't confuse URL filtering with application detection, which relies on reading some of the packet from a website to determine more specifically what it is (for example, Facebook Message or Salesforce). For more information, see Recommendations for application control, on page 2054.

QUIC Fingerprinting

Snort can identify client applications in Quick UDP Internet Connections (QUIC sessions) based on EVE analysis. Refer to QUIC fingerprinting in EVE, on page 2241 for more information.

License requirements for URL filtering

This reference provides the license requirements for URL filtering functionality on Cisco Secure Firewall devices, covering both category and reputation filtering as well as manual filtering options.

Threat Defense License

- Category and reputation filtering—URL Filtering
- Manual filtering—No additional license.

Classic License

- Category and reputation filtering—URL Filtering

- Manual filtering—No additional license.

URL filtering licenses for threat defense devices

Refer to *URL Licenses* in the *Licenses* chapter in the Cisco Secure Firewall Management Center Administration Guide.

Requirements and prerequisites for URL filtering

This reference provides the requirements and prerequisites necessary for implementing URL filtering, including model support, supported domains, and required user roles.

Model support

Any

Secure Firewall 200 does not maintain a local URL database. These options are not relevant for this device type and are automatically changed to Cloud Only lookup when deploying policies: **Enable Automatic Updates**, URL Query Source for the **Local Database Only** and the **Local Database and Cisco Cloud** options.

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Configure URL filtering with category and reputation

This task enables URL filtering with category and reputation to protect your network from malicious websites and to control user access to specific website categories.

URL filtering with category and reputation provides comprehensive web security by evaluating websites based on their content categories and reputation scores. This feature helps prevent access to malicious sites while allowing administrators to control access to specific types of content.

Before you begin

- Assign the URL Filtering license to each managed device that will filter URLs.

Follow these steps to configure URL filtering with category and reputation:

Procedure

Step 1 Review and understand limitations AND guidelines.

Best practices for URL filtering, on page 2088

Step 2 Enable the URL Filtering feature.

Enable URL filtering using category and reputation, on page 2096

Step 3 Configure rules to filter URLs by category and reputation.

Configure URL conditions, on page 2097

For the best protection against malicious sites, you must block sites by reputation and block URLs in all threat categories.

(Optional) Supplement or selectively override category and reputation-based URL filtering, on page 2102

Step 4 (Optional) Allow users to bypass a website block by clicking through a warning page.

This step is optional.

HTTP response pages, on page 2103

Step 5 Order your rules so that traffic hits the key rules first.

URL rule order, on page 2099

Step 6 (Optional) Modify advanced options related to URL filtering.

This step is optional. Generally, use the defaults unless you have a specific reason to change them.

For information about advanced options, refer to Advanced Settings for Access Control, on page 2687.

- **Maximum URL characters to store in connection events**
- **Allow an Interactive Block to bypass blocking for (seconds)**
- **Retry URL cache miss lookup**
- **Enable reputation enforcement on DNS traffic**

Step 7 Deploy your changes.

Deploy configuration changes, on page 202

Step 8 Ensure that your system receives future URL data updates as expected.

Configure URL filtering health monitors, on page 2107

Step 9 Enable other features to protect your network from malicious sites.

See Security Intelligence, on page 2027.

Enable URL filtering using category and reputation

Enable URL filtering to control and monitor web traffic based on website categories and reputation scores, providing enhanced security for your network.

You must be an Admin user to perform this task.

Before you begin

Complete prerequisites described in Configure URL filtering with category and reputation, on page 2094.

Follow these steps to enable URL filtering using category and reputation:

Procedure

Step 1 Choose **Integration** > **Other Integrations** > **Cloud Services**.

Step 2 Configure URL filtering options, on page 2096.

Step 3 Click **Save**.

URL filtering options

Adding a URL Filtering license automatically enables the URL filtering feature. With this feature, the system handles traffic based on a website's general classification (category) and risk level (reputation).

Although by default the system is configured to submit all URLs to the cloud for threat intelligence evaluation, using a local dataset of category and reputation data can make web browsing faster.When you enable or re-enable URL filtering, the Cloud-Delivered Firewall Management Center automatically queries Cisco for URL data and pushes the dataset to managed devices. This process may take some time.

If you use SSL rules to handle encrypted traffic, also see Rule-based decryption rules guidelines and limitations, on page 2188.

Enable automatic updates

If you **Enable Automatic Updates** (the default), the Cloud-Delivered Firewall Management Center checks the cloud every 30 minutes for updates. If you need strict control over when the system contacts external resources, disable automatic updates and instead create a recurring task using the scheduler. See the *Cisco Secure Firewall Management Center Administration Guide*.

Update now

Click **Update Now** to perform a one-time, on demand URL data update. You cannot start an on demand update if an update is already in progress. Although daily updates tend to be small, if it has been more than five days since your last update, new URL data may take as long as 20 minutes to download, depending on your bandwidth. The update itself may take up to 30 minutes to complete.

URL query source

You can choose how the system assigns a category and reputation to the URLs that your users browse to. You can choose:

- **Local Database Only**: Uses the local dataset only. Use this option if you do not want to submit your uncategorized URLs (category and reputation not in the local dataset) to Cisco, for example, for privacy reasons. However, note that connections to uncategorized URLs do *not* match rules with category or reputation-based URL conditions. You cannot assign categories or reputations to URLs manually.
- **Local Database and Cisco Cloud**: Uses the local dataset when possible, which can make web browsing faster. When users browse to an URL whose category and reputation is not in the local dataset or a cache of previously accessed websites, the system submits it to the cloud for threat intelligence evaluation and adds the result to the cache.
- **Cisco Cloud Only** (default): Does not use the local dataset. When users browse to an URL whose category and reputation is not in a local cache of previously accessed websites, the system submits it to the cloud for threat intelligence evaluation and adds the result to the cache. This option guarantees the most up-to-date category and reputation information.

  This option requires Version 7.3 on the device. If you enable this option, devices running earlier versions will use the **Local Database and Cisco Cloud** option.

Cached URLs expire

Caching category and reputation data makes web browsing faster. By default, cached data for URLs never expires, for fastest performance.

To minimize instances of URLs matching on stale data, you can set cached URLs to expire. For more accurate and current threat data, choose a shorter expiration time. A cached URL refreshes after the first time a user on the network visits it after the specified time passes. The first user does not see the refreshed result, but the next user who visits the URL does.

Configure URL conditions

Configure URL conditions to protect your network by controlling access to sites based on URL category and reputation.

URL conditions allow you to control network access based on website categories and reputations, providing protection against malicious sites and enabling security policy enforcement.

Before you begin



Attention As a prerequisite, ensure that you create at least a Monitor rule at the top of your access control policy, containing Category or Reputation parameters. This is essential to see ANY category or reputation data for ANY URLs that hit the particular access control policy.

If there is no rule in the access control policy with the category or reputation parameters configured, the **Connection Events** page in the management center shows no data for Category or Reputation for ANY URL traffic that hits the access control policy.

Follow these steps to configure URL conditions:

Procedure

Step 1 In the rule editor, click the following for URL conditions:

- Access control or QoS—Click **URLs**.
- SSL—Click **Category**.

Step 2 Find and choose the URL categories that you want to control:

In an access control or QoS rule, click **Category**.

For effective protection from malicious sites, you must block URLs in all Threat categories. Additionally, Talos recommends that you block only sites with Poor category. You can block questionable reputations if you have an aggressive security posture, but this may result in a higher amount of false positives. For a list of Threat categories, see URL category and reputation descriptions, on page 2087.

Be sure to click the arrows at the bottom of the list to see all available categories.

Step 3 (Optional) Constrain URL categories by choosing a **Reputation**.

Note that if you explicitly match **Uncategorized** URLs, you cannot further constrain by reputation. Choosing a reputation level also includes other reputations either more or less severe than the level you choose, depending on the rule action:

- Includes less severe reputations—If the rule allows or trusts web traffic. For example, if you configure an access control rule to allow Favorable (level 4), it also automatically allows Trusted (level 5) sites.
- Includes more severe reputations—If the rule rate limits, decrypts, blocks, or monitors web traffic. For example, if you configure an access control rule to block Questionable sites (level 2), it also blocks Untrusted (level 1) sites.

If you change the rule action, the system automatically changes the reputation levels in URL conditions.

Optionally, select **Apply to unknown reputation**.

Step 4 Click **Add URL** or **Add to Rule**, or drag and drop.

Step 5 (Optional) To choose predefined URL objects, or URL lists and feeds in an access control or QoS rule, click **URL**, select the objects, and add them to the destination.

These objects implement manual URL filtering rather than category-based filtering.

Step 6 Save or continue editing the rule.

URL condition in an access control rule

The following graphic shows the URL condition for an access control rule that blocks all malware sites, all untrusted sites, and all social networking sites with a reputation level of Neutral or worse.



This table summarizes how you build the condition.

Blocked URL	Category	Reputation
Malware sites, regardless of reputation	Malware Sites	ANY
Any untrusted URL (level 1)	ANY	1 - Untrusted
Social networking sites with a reputation level of Neutral or worse (levels 1 through 3)	Social Network	3 - Neutral

Rules with URL conditions

This table lists rules that support URL conditions, and the types of filtering that each rule type supports.

Rule Type	Supports Category and Reputation Filtering?	Supports Manual Filtering?
Access control	Yes	Yes
Decryption policy	Yes	No. Use distinguished name conditions instead.
QoS	Yes	Yes

To use URL filtering in a decryption policy that has **Do Not Decrypt** rule conditions, refer to Use categories in URL filtering, on page 2092.

URL rule order

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.
- If you configure exceptions to a rule, place the exception above the other rule.

DNS filtering

DNS filtering is a URL filtering feature that

- evaluates domain category and reputation early in URL transactions, when the browser looks up the domain name to get the IP address,
- can often determine category and reputation of encrypted traffic without decryption, and
- is enabled by default through the **Enable reputation enforcement on DNS traffic** option on the **Advanced** tab of each new access control policy.

DNS filtering behavior

DNS filtering slightly modifies URL filtering behavior and is applicable only when URL filtering is enabled and configured.

If DNS filtering cannot determine the URL of encrypted traffic, that traffic is processed using your configurations for encrypted traffic.

Enable DNS filtering to identify URLs during domain lookup

DNS filtering allows you to identify and block URLs based on category and reputation during domain lookup, providing an additional layer of security by preventing access to malicious or unwanted domains before connections are established.

DNS filtering is enabled by default in new access control policies. However, additional configurations may be required in order for this setting to take effect.

Before you begin

- URL filtering using category and reputation must be licensed, enabled, and configured.

  DNS filtering does not use these settings in the URLs tab: URL groups, URL objects, URL lists and feeds, and URLs entered into the "Enter URL" text box.
- Refer to limitations at DNS filtering behavior for block actions, on page 2101.

Procedure

Step 1 In your access control policy's advanced settings, select **Enable reputation enforcement on DNS traffic**.

Step 2 In the same policy, for each access control rule that has URL category and reputation blocking configured:

- Application conditions—If the application condition is anything other than **any** (or empty), add **DNS** to that list. Other DNS-related options are not relevant for this purpose.
- Port condition—If the port/protocol condition is anything other than **any** (or empty), add **DNS_over_TCP** and **DNS_over_UDP**.

Step 3 Save your changes.

What to do next

If you are done making changes, Deploy configuration changes, on page 202.

DNS filtering behavior for block actions

Traffic that matches rules having action **Block with reset**, **Interactive Block**, or **Interactive Block with reset** will be treated as if the rule action were **Block**.

End users trying to access a blocked URL will experience this as an unexplained inability to connect to their page. The connection will spin and then time out.

DNS filtering and events

DNS filtering and events is a logging mechanism that

- generates connection events using specific fields such as DNS Query, URL Category, URL Reputation, and Destination Port,
- populates the DNS Query field with the domain name while leaving the URL field blank for DNS filtering matches, and
- sets the Destination Port to 53.

Event generation behavior

DNS filtering generates events based on the access control rule action:

- When the access control rule action is **Allow** or **Trust**, two connection events will be generated for the same traffic, one for DNS filtering (with the **DNS Query** field populated) and one for URL filtering (with the **URL** field populated).
- The first time the system encounters a particular URL, two events appear for that single session: one event showing uncategorized or reputationless for the DNS Query, and one event showing the actual category and reputation for the URL, which the system retrieves during the DNS Query and applies to the session while processing using standard URL filtering.

Manual URL filtering

Manual URL filtering is a URL filtering capability that

- supplements or selectively overrides category and reputation-based URL filtering by manually filtering individual URLs, groups of URLs, or URL lists and feeds in access control and QoS rules, and
- enables URL filtering without requiring a special license when manually permitting or blocking URLs through Allow or Block rules within an access control policy.

Manual URL filtering implementation examples

Examples:

- You might use access control to block a category of websites that are not appropriate for your organization. However, if the category contains a website that is appropriate, and to which you want to provide access, you can create a manual Allow rule for that site and place it before the Block rule for the category.
- If you prefer not to use category-based URL filtering, you can manually permit or block a URL or a group of URLs by creating an Allow or Block rule within an access control policy. You can perform this type of URL filtering without a special license.

Manual URL filtering is not supported in SSL rules. Use distinguished name conditions instead.

Caution Depending on how you implement manual URL filtering, URL matching may not be what you intend. Refer to Manual URL filtering options, on page 2102.

Manual URL filtering options

There are several ways to specify URLs for manual URL filtering.

Option	Description
(Best practice) Use custom Security Intelligence URL list or feed objects.	This is the recommended method for manual URL filtering. You can create a new list or feed, or choose an existing one in an access control or QoS rule. For more information, see Custom security intelligence lists and feeds, on page 3026 and subtopics.
Use URL objects, individually or as groups. URL objects are described at URLs, on page 3038. Or Enter URLs directly into the access control rule. (The **Enter URL** option on the rule page in the web interface.)	If you do not include a path (that is, there are no slash (/) characters in the URL), the match is based on the server's hostname only. If you include one or more / character, the entire URL string is used for a substring match. Then, a URL is considered a match if any of these are true: • The string is at the beginning of the URL. • The string follows a dot. • The string contains a dot in the beginning. • The string follows the :// characters. For example, ign.com matches ign.com or www.ign.com, but not versign.com. **Note** We recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites (that is, URL strings with / characters), as servers can be reorganized and pages moved to new paths. The **Enter URL** option does not support wildcards.

Supplement or selectively override category and reputation-based URL filtering

This task allows you to use Security Intelligence URL lists and feeds to supplement or create exceptions to your category and reputation-based URL filtering rules in access control or QoS policies.

In access control or QoS rules, you can use Security Intelligence URL lists and feeds to supplement, or to specify exceptions to, your category and reputation-based URL filtering rules.

Important! If the list or feed you are configuring in this procedure contains exceptions to category- or reputation-based rules, put this rule above those rules in the rule order.

In SSL rules, use distinguished name conditions to configure parallel behavior.

Before you begin

- Configure URL filtering using category and reputation. Refer to Configure URL conditions, on page 2097.
- Understand important best practices for manual URL filtering. Refer to Best practices for URL filtering, on page 2088 and Manual URL filtering options, on page 2102.
- Configure one or more Security Intelligence objects (lists or feeds) containing the URLs that you want to use for manual filtering. Refer to Custom security intelligence lists and feeds, on page 3026.

Follow these steps to supplement or selectively override category and reputation-based URL filtering:

Procedure

Step 1 Navigate to the access control or QoS policy in which you will define the rule.

Step 2 Create or edit the rule in which you will add the new condition:

- If you are supplementing a category- or reputation-based URL filtering rule, edit the existing rule.
- If you are overriding or creating exceptions to a category- or reputation-based URL filtering rule, create a new rule.

Step 3 Select the list or feed you created as the destination URL criteria.

Step 4 Save the rule.

HTTP response pages

An HTTP response page is a web page that displays when the system blocks web requests using either access control rules or the access control policy default action.

HTTP response page types

The response page displayed depends on how you block the session:

- **Block Response Page**: Overrides the default browser or server page that explains that the connection was denied.
- **Interactive Block Response Page**: Warns users, but also allows them to click a button (or refresh the page) to load the originally requested site. Users may have to refresh after bypassing the response page to load page elements that did not load.

If you do not choose a response page, the system blocks sessions without interaction or explanation.

Limitations to HTTP response pages

The system displays a response page only for unencrypted or decrypted HTTP/HTTPS connections blocked (or interactively blocked) either by access control rules or by the access control policy default action. The system does not display a response page for connections that are blocked by any other policy or mechanism.

The system cannot display a response page if the connection is reset (RST packet sent). If you enable response pages, the system prioritizes that configuration. Even if you choose **Block with reset** or **Interactive Block with reset** as the rule action, the system displays the response page and does not reset matching web connections. To ensure that blocked web connections are reset, you must disable response pages.

- Note that all non-web traffic that matches the rule *is* blocked with reset.

The system does not display a response page for encrypted connections that are blocked by access control rules (or any other configuration). Access control rules evaluate encrypted connections if you did not configure an SSL policy, or your SSL policy passes encrypted traffic.

-
- However, the system does display a response page for connections that are decrypted by the SSL policy, then blocked (or interactively blocked) either by access control rules or by the access control policy default action. In these cases, the system encrypts the response page and sends it at the end of the reencrypted SSL stream.

The system does not display a response page when web traffic is blocked because of a promoted access control rule (an early-placed blocking rule with only simple network conditions).

If a URL is entered without specifying "http" or "https", and the browser initiates the connection on port 80, and the user clicks through a response page, and the connection is subsequently redirected to port 443, the user will not see a second interactive response page because the response to this URL is already cached.

The system does not display a response page when web traffic is blocked before the system identifies the requested URL. Refer to Best practices for URL filtering, on page 2088.

The system does not display a response page if the block URL access control rule is configured after the allow application rule.

Requirements and prerequisites for HTTP response pages

Model support

Any

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Describe the necessary conditions, user roles, and supported domains required for setting up HTTP response pages.

Choose HTTP response pages

HTTP response pages provide feedback to users when their sessions are blocked by access control policies, allowing for better user experience and network transparency.

Reliable display of HTTP response pages depends on your network configuration, traffic loads, and size of the page. Smaller pages are more likely to display successfully.

Procedure

Step 1 In the access control policy editor, select **HTTP Responses** from the **More** drop-down arrow at the end of the packet flow line.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Choose the **Block Response Page** and **Interactive Block Response Page**:

- System-provided—Displays a generic response. Click **View** ( ) to view the code for this page.
- Custom—Create a custom response page. A pop-up window appears, prepopulated with system-provided code that you can replace or modify by clicking **Edit** ( ). A counter shows how many characters you have used.
- None—Disables the response page and blocks sessions without interaction or explanation. To quickly disable interactive blocking for the whole access control policy, choose this option.

Step 3 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Interactive blocking with HTTP response pages

Interactive blocking with HTTP response pages is a security mechanism that

- allows users to load an originally requested site after reading a warning page,
- provides user control over whether to proceed with potentially blocked traffic, and
- enables temporary bypass with configurable duration settings.

Configuration behavior and bypass settings

When you configure interactive blocking, users may have to refresh after bypassing the response page to load page elements that did not load.

Tip To quickly disable interactive blocking for the whole access control policy, display neither the system provided page nor a custom page. The system then blocks all connections without interaction.

If a user does not bypass an interactive block, matching traffic is denied without further inspection. If a user bypasses an interactive block, the access control rule allows the traffic, although the traffic may still be subject to deep inspection and blocking.

By default, a user bypass is in effect for 10 minutes (600 seconds) without displaying the warning page on subsequent visits. You can set the duration to as long as a year, or you can force the user to bypass the block every time. This limit applies to every Interactive Block rule in the policy. You cannot set the limit per rule.

Logging options for interactively blocked traffic are identical to those in allowed traffic. If a user does not bypass the interactive block, the system can log only beginning-of-connection events. When the system initially warns the user, it marks any logged beginning-of-connection event with the `Interactive Block` or `Interactive Block with reset` action. If the user bypasses the block, additional connection events logged for the session have an action of `Allow`.

Configure interactive blocking

Interactive blocking allows users to bypass URL filtering rules while maintaining security controls and logging capabilities.

This procedure describes how to allow users to bypass URL filtering rules.

Procedure

Step 1 As part of access control, configure an access control rule that matches web traffic; see Create and edit access control rules, on page 2066:

- Action—Set the rule action to **Interactive Block** or **Interactive Block with reset**; see Access control rule interactive blocking actions, on page 2047.
- Conditions—Use URL conditions to specify the web traffic to interactively block.
- Logging—Assume users will bypass the block and choose logging options accordingly.
- Inspection—Assume users will bypass the block and choose deep inspection options accordingly; see Access control policies, on page 1997.

Step 2 (Optional) (Optional) In the access control policy **HTTP Responses**, choose a custom interactive-block HTTP response page; see Choose HTTP response pages, on page 2105.

Step 3 (Optional) (Optional) In the access control policy **Advanced** settings, change the user bypass timeout; see Set the user bypass timeout for a blocked website, on page 2107.

After a user bypasses a block, the system allows the user to browse to that page without warning until the timeout period elapses.

Step 4 Save the access control policy.

Step 5 Deploy configuration changes.

Set the user bypass timeout for a blocked website

Set the duration that users can browse after bypassing a URL filtering block before having to bypass the block again.

This procedure explains how to set the time allowed for browsing after the user bypasses a URL filtering block. When the timeout expires, the user must bypass the block again.

Procedure

Step 1 Click **Policies** > **Access Control heading** > **Access Control** and edit the policy.

Step 2 Select **Advanced Settings** from the **More** drop-down menu at the end of the packet flow line.

Step 3 Click **Edit** ( ) next to General Settings.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 In the **Allow an Interactive Block to bypass blocking for (seconds)** field, enter the number of seconds that must elapse before the user bypass expires.

Setting this value to 0 means the interactive block response is displayed once and the user bypass never expires.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Configure URL filtering health monitors

The following health policies alert if the system has problems obtaining or updating URL category and reputation data.

- URL Filtering Monitor
- Threat Data Updates on Device

To ensure that these are configured the way you want them, refer to *Health Modules* and *Configuring Health Monitoring* in the *Cisco Secure Firewall Management Center Administration Guide*.

Dispute URL category and reputation

This task allows you to challenge URL categorizations or reputation scores when you disagree with the assignments made by Talos Intelligence.

If you disagree with a category or reputation assigned by Talos, you can submit a request for re-evaluation.

Before you begin

You will need your Cisco account credentials.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center web interface, do one of the following:

Location of dispute option	Path to dispute option
Cloud Services configuration page	a. Navigate to the **Integration** > **Other Integrations** > **Cloud Services** page. b. Select **Dispute URL categories and reputations**.

The Talos web site opens in a separate browser window.

Step 2 Sign in to the Talos site with your Cisco credentials.

Step 3 Review the information and follow the instructions on the Talos page.

Step 4 On the Talos site, review how submitted disputes are handled. Check what response you can expect, if any.

The dispute process is independent of Secure Firewall products.

Take action if the URL category set changes

URL filtering categories may occasionally change to accommodate new web trends and evolving usage patterns. Taking action when these changes occur ensures that your policies continue to function properly and can be deployed successfully.

The set of URL Filtering categories may occasionally change, in order to accommodate new web trends and evolving usage patterns.

These changes affect both policies and events.

Shortly before URL category changes are scheduled to occur, and after they occur, you will see alerts in the list of rules in any access control, SSL, and QoS policy that is affected by the changes, and on URL or Category in rules that you edit.

You should take action when you see these alerts.

Note Updates to the URL category set as described in this topic are distinct from the changes that simply add new URLs to existing categories or re-classify misclassified URLs. This topic does not apply to category changes for individual URLs.

Procedure

Step 1 If you see an alert beside a rule in an access control policy, hover over the alert to see details.

Step 2 If the alert mentions changes to URL categories, edit the rule to see further details.

Step 3 Hover over the URL or Category in the rule dialog to see general information about the type of changes.

Step 4 If you see an alert beside a category, click the alert to view details.

Step 5 If you see a "More information" link in the description of a change, click it to view information about the category on the Talos web site.

Alternatively, see a list and descriptions of all categories at the link in URL category and reputation descriptions, on page 2087.

Step 6 Depending on the type of change, take appropriate action:

Type of Category Change	What The System Will Do	What You Should Do
Existing category will soon be deprecated	Nothing yet. You have a few weeks to change affected rules. If you do not take action in that time, the system eventually will not be able to redeploy the policy.	Remove this category from all rules that include it. If there is a similar new category, consider using that category instead.
New category is added	By default, the system does not use newly added categories.	Consider creating new rules for the new category.
Existing category is deleted	The category will appear in the rule in strikethrough text (that is, with a line through the category name.)	You must delete the obsolete category from the rule before you can deploy the policy.

Step 7 Check your SSL rules (Category) for these changes and take action as needed.

Step 8 Check your QoS rules (URL) for these changes and take action as needed.

What to do next

Deploy configuration changes.

URL category and reputation changes

URL category and reputation changes are system modifications that

- affect how events are categorized and labeled in the system,
- result in different handling of events processed before and after the changes, and
- impact the display of URL reputation information in event viewers.

Effects on event processing

URL category and reputation changes affect events differently based on when they were processed:

- When URL categories change, events that the system processed before the category change will be associated with their original category names and will be labeled with **Legacy**. Events that the system processed after the category change will be associated with the new categories.

  Older, legacy events will age out of the system over time.

- If a URL does not have a reputation at the time it was processed, the URL Reputation column in the event viewer will be empty.

Troubleshoot URL filtering

This reference provides troubleshooting information for common URL filtering issues including missing categories, uninspected packets, registration failures, lookup errors, and performance problems.

Expected URL category is missing from the categories list

The URL filtering feature uses a different set of categories than the Security Intelligence feature. The category that you expect to see may be a Security Intelligence category. To see those categories, look at the **URLs** tab on the **Security Intelligence** tab in an access control policy.

Initial packets are passing uninspected

Refer to Inspection of packets that pass before traffic is identified, on page 2918 and subtopics.

See also DNS filtering, on page 2099.

How can I find the category and reputation of a particular URL?

Do a manual lookup. Refer to *Finding URL Category and Reputation* in the *Cisco Secure Firewall Management Center Administration Guide*.

Error when attempting a manual lookup: "cloud lookup failure for <URL>"

Make sure the feature is properly enabled. Refer to the prerequisites in *Finding URL Category and Reputation* in the *Cisco Secure Firewall Management Center Administration Guide*.

URL appears to be incorrectly handled based on its URL category and reputation

Problem: The system does not handle the URL correctly based on its URL category and reputation.

Solutions:

- Verify that the URL category and reputation associated with the URL are what you think they are. See *Finding URL Category and Reputation* in the *Cisco Secure Firewall Management Center Administration Guide*.
- These issues may be addressed by settings described in URL filtering options, on page 2096, accessible using Enable URL filtering using category and reputation, on page 2096.
  - The URL cache may hold stale information. See information about the **Cached URLs Expire** setting in URL filtering options, on page 2096.

  - The local data set may not be updated with current information from the cloud. See information about the **Enable Automatic Updates** setting in URL filtering options, on page 2096.
  - The system may be configured to *not* check the cloud for current data. See information about the **Query Cisco cloud for unknown URLs** setting in URL filtering options, on page 2096.

- Your access control policy may be configured to pass traffic to the URL without checking the cloud. See information about the **Retry URL cache miss lookup** setting in General settings, on page 2691.
- See also Best practices for URL filtering, on page 2088.
- If the URL is processed using an SSL rule, see Rule-based decryption rules guidelines and limitations, on page 2188 and Best practice for Decryption policy rule order, on page 2196
- Verify that the URL is being handled using the access control rule that you think it is being handled by, and that the rule does what you think it does. Consider rule order.
- Verify that the local URL category and reputation database on the Cloud-Delivered Firewall Management Center is successfully being updated from the cloud and that managed devices are successfully being updated from the Cloud-Delivered Firewall Management Center.

  Status of these processes are reported in the Health Monitor, in the **URL Filtering Monitor** module and the **Threat Data Updates on Devices** module. For details, see *Health* in the *Cisco Secure Firewall Management Center Administration Guide*.

If you want to immediately update the local URL category and reputation database, go to **Integration** > **Other Integrations** > **Cloud Services**, then click **Update Now** in the **URL Filtering** area. For more information, see URL filtering options, on page 2096.

A URL category or reputation is not correct

For access control or QoS rules: Use manual filtering, paying careful attention to rule order. See Manual URL filtering, on page 2101 and Configure URL conditions, on page 2097.

For SSL rules: Manual filtering is not supported. Instead, use distinguished name conditions.

See also Dispute URL category and reputation, on page 2107.

Web pages are slow to load

There is a tradeoff between security and performance. Some options:

- Consider modifying the **Cached URLs Expire** setting. Choose **Integration** > **Other Integrations** > **Cloud Services**. For information, see URL filtering options, on page 2096.
- Consider deselecting the **Retry URL cache miss lookup** setting in General settings, on page 2691.

Events do not include URL category and reputation

- Make sure you have included applicable URL rules in an access control policy, the rules are active, and the policies have been deployed to the relevant devices.
- URL category and reputation do not appear in an event if the connection is processed before it matches a URL rule.
- The rule that handles the connection must be configured for URL category and reputation.

- Even if you have configured URL categories in the Categories tab in an SSL rule, you must also configure the URLs tab in a rule in your access control policy.

DNS filtering is not working

Make sure you have completed all prerequisites and steps in Enable DNS filtering to identify URLs during domain lookup, on page 2100.

An end user tries to access a blocked URL and the page just spins and times out

When DNS Filtering is enabled and end users access a URL that is blocked, the page will spin but not load. End users are not notified that the page is blocked. This is currently a limitation when DNS filtering is enabled.

Refer to DNS filtering behavior for block actions, on page 2101.

Events include URL category and reputation but URL field is blank

If the DNS Query field is populated and the URL field is empty, this is expected when the DNS filtering feature is enabled.

Refer to DNS filtering and events, on page 2101.

Multiple events are generated for a single transaction

A single web transaction sometimes generates two connection events, one for DNS filtering and one for URL filtering. This is expected when DNS filtering is enabled and:

- the access control rule action for the traffic is Allow or Trust.
- the system encounters a URL for the first time.

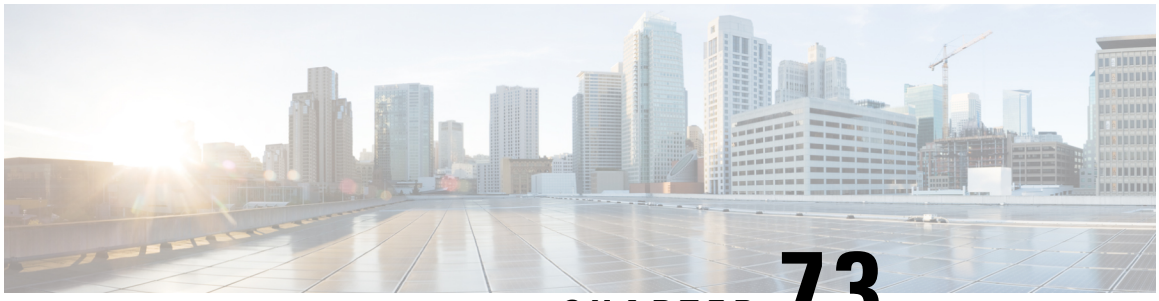
Refer to DNS filtering and events, on page 2101.



PART XVIII

Decryption Policies and Encrypted Visibility for Access Control

- Traffic Decryption Overview, on page 2115
- Standard Decryption Policies, on page 2137
- Rule-Based Decryption Policies, on page 2163
- Rule-Based Decryption Rules, on page 2187
- Encrypted Visibility Engine, on page 2239



CHAPTER 73

Traffic Decryption Overview

These topics provide an overview of Transport Layer Security/Secure Sockets Layer (TLS/SSL) inspection, discuss the prerequisites for TLS/SSL inspection configuration, and detail deployment scenarios.

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

For more information about SSL and TLS protocols, see a resource such as SSL vs. TLS - What's the Difference?

- Traffic decryption, on page 2115
- How TLS/SSL handshakes work, on page 2116
- Decryption rule and policy basics, on page 2121
- TLS crypto acceleration, on page 2130
- History for decryption policy , on page 2133

Traffic decryption

Traffic decryption enables deep inspection of traffic for analysis.

Traffic handling options

Most traffic on the internet is encrypted and in most cases, you might choose not to decrypt it. You can glean some information from encrypted traffic and you can also block it from your network if desired.

Your choices are:

- Decrypt the traffic and subject it to the full array of deep inspection:
  - Advanced Malware Protection
  - Security intelligence
  - Threat Intelligence Director
  - Application detectors

- URL and category filtering

- Leave the traffic encrypted and set up your access control and decryption policy to look for and potentially block:
  - Old protocol versions (such as Secure Sockets Layer)
  - Unsecure cipher suites
  - Applications with high risk and low business relevance
  - Untrusted issuer Distinguished Names

Access control policy integration

An access control policy is the main configuration that invokes subpolicies and other configurations, including a decryption policy. If you associate a decryption policy with access control, the system uses that decryption policy to handle encrypted sessions before it evaluates the sessions with access control rules. If you do not configure TLS/SSL inspection, or your devices do not support it, access control rules handle all encrypted traffic.

Access control rules also handle encrypted traffic when your TLS/SSL inspection configuration allows the traffic to pass. However, some access control rule conditions require unencrypted traffic, so encrypted traffic might match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improves performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Note As long as your managed devices have Snort 3 enabled, the system supports decrypting TLS 1.3 traffic. You can enable TLS 1.3 decryption in an decryption policy's advanced options; for more information, see Rule-based decryption policy advanced options, on page 2182.

Notes

Decryption rules require processing overhead that can impact performance. Before you decide to decrypt traffic, see Best practice for deciding when to decrypt traffic, on page 2122.

The system does not support mutual authentication; that is, you cannot upload a client certificate to the Cloud-Delivered Firewall Management Center and use it for **Decrypt - Resign**, **Decrypt - Replace Cert**, or **Decrypt - Known Key** decryption rule actions.

If you set the value of TCP maximum segment size (MSS) using FlexConfig, the observed MSS could be less than your setting. For more information, see TCP MSS, on page 1106.

How TLS/SSL handshakes work

In this documentation, the term *TLS/SSL handshake* represents the two-way handshake that initiates encrypted sessions in both the SSL protocol and its successor protocol, TLS.

In an inline deployment, the system processes the TLS/SSL handshake, potentially modifying the ClientHello message and acting as a TCP proxy server for the session.

The figure shows an inline deployment.



Summary

The key components involved in the TLS/SSL handshake process are:

- Client: Initiates the encrypted session by sending a ClientHello message specifying multiple supported values for each encryption attribute
- Server: Responds with a ServerHello message specifying a single supported value for each encryption attribute, determining the encryption method for the secure session
- Managed device: Monitors TCP sessions for encrypted session initiation attempts and can act as a TCP proxy server
- Handshake messages: Specialized packets that communicate encryption attributes between client and server

Workflow

These stages describe how TLS/SSL handshakes establish encrypted sessions:

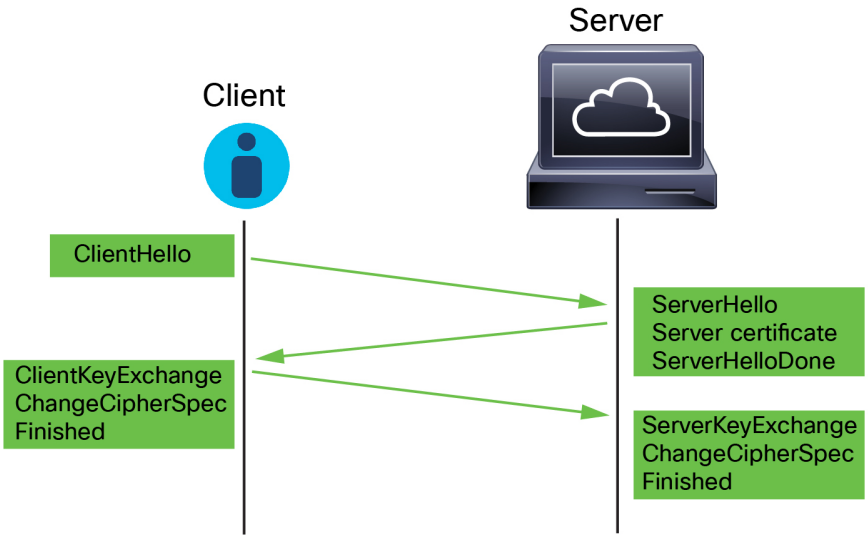
1. The client establishes a TCP connection with the server after successfully completing the TCP three-way handshake.
2. The managed device monitors the TCP session for any attempt to initiate an encrypted session.
3. The client and server exchange specialized handshake messages to establish an encrypted session.
   - ClientHello—The client specifies multiple supported values for each encryption attribute.
   - ServerHello—The server specifies a single supported value for each encryption attribute, and the ServerHello response determines which encryption method the system uses during the secure session.
4. After the TLS/SSL handshake completes, the managed device caches encrypted session data and server certificate data.

   This caching allows session resumption without requiring the full handshake and enables faster handshake processing in subsequent sessions that use the same certificate.

ClientHello message handling

The client sends the ClientHello message to the server, which acts as the packet destination, when a secure connection can be established. The client uses this message to initiate the TLS/SSL handshake to respond to a ServerHello message from the destination server.

The following figure shows an example. Also see RFC 8446, sec. 4 . You can also consult a resource such as What Happens in a TLS Handshake? at cloudflare.com.



ClientHello message process

The process consists of these steps:

1. ClientHello initiates the process.

   The ClientHello message contains the Server Name Indication (SNI), which has the server's fully qualified domain name.

2. After a managed device processes a ClientHello message and transmits it to the destination server, the server determines whether it supports the decryption attributes the client specified in the message. If it does not support those attributes, the server sends a handshake failure alert to the client. If it supports those attributes, the server sends the ServerHello message. If the agreed-upon key exchange method uses certificates for authentication, the server certificate message immediately follows the ServerHello message.

   The server certificate contains the Subject Alternative Name (SAN), which can have fully qualified domain names and IP addresses. For more information about the SAN, see Distinguished names, on page 2962.

3. When the managed device receives these messages, it attempts to match them with decryption rules configured on the system. These messages contain information that was absent from either the ClientHello message or the session data cache. Specifically, the system can potentially match these messages on decryption rules' Distinguished Names, Certificate Status, Cipher Suites, and Versions conditions.

The entire process is encrypted.

Data exchange

If you configure TLS/SSL decryption, when a managed device receives a ClientHello message, the system attempts to match the message to decryption rules that have the **Decrypt - Resign** , or **Decrypt - Known Key** action. The match relies on data from the ClientHello message and from cached server certificate data. The table lists the possible types of data available for matching conditions.

Table 122: Data Availability for Decryption rule Conditions

Decryption rule Condition	Data Present In
Zones	ClientHello
Networks	ClientHello
VLAN Tags	ClientHello
Ports	ClientHello
Users	ClientHello
Applications	ClientHello (Server Name Indicator extension)
Categories	ClientHello (Server Name Indicator extension)
Certificate	Server certificate (potentially cached)
Distinguished Names	Server certificate (potentially cached)
Certificate Status	Server certificate (potentially cached)
Cipher Suites	ServerHello
Versions	ServerHello



Important Use the **Cipher Suite** and **Version** rule conditions *only* in rules with either the **Block** or **Block with reset** rule actions. Do not use **Cipher Suite** and **Version** with **Decrypt - Resign** or **Decrypt - Known Key** rule actions. These conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

ClientHello message modifications

If the ClientHello message matches a **Decrypt - Resign** , or **Decrypt - Known Key** rule, the system modifies the ClientHello message.

- (TLS 1.2 only; TLS 1.3 does not support compression.) Compression methods—Strips the `compression_methods` element, which specifies the compression methods the client supports. The system cannot decrypt compressed sessions.
- Cipher suites—Strips cipher suites from the `cipher_suites` element if the system does not support them. If the system does not support any of the specified cipher suites, the system transmits the original, unmodified element. This modification reduces the Unknown Cipher Suite and Unsupported Cipher Suite types of undecryptable traffic.
- Session identifiers—Strips any value from the `Session Identifier` element and the SessionTicket extension (RFC 5077, sec 3.2) that does not match cached session data. If a ClientHello value matches cached data, an interrupted session can resume without the client and server performing the full TLS/SSL handshake. This modification improves the likelihood of session resumption and reduces the Session Not Cached incidents of undecryptable traffic..

- Elliptic curves—Strips elliptic curves from the Supported Elliptic Curves extension if the system does not support them. If the system does not support any of the specified elliptic curves, the managed device removes the extension and strips any related cipher suites from the `cipher_suites` element.
- ALPN extensions—Strips any value from the Application-Layer Protocol Negotiation (ALPN) extension that is unsupported in the system (for example, the HTTP/2 protocol).
- Other Extensions—Strips the Next Protocol Negotiation (NPN) and TLS Channel IDs extensions.

  Decryption rules with a **Decrypt - Resign** , or **Decrypt - Known Key** action now natively support the Extended Master Secret (EMS) extension during ClientHello negotiation, enabling more secure communications. The EMS extension is defined by RFC 7627 .

After the system modifies the ClientHello message, it determines whether the message passes access control evaluation (which can include deep inspection). If the message passes, the system transmits it to the destination server.

If the ClientHello message does *not* match a **Decrypt - Resign** , or **Decrypt - Known Key** rule, the system does not modify the message. It then determines whether the message passes access control evaluation (which can include deep inspection). If the message passes inspection, the system transmits it to the destination server.

ClientHello is *not* modified if traffic matches a **Monitor** rule condition.

Man-in-the-middle

Direct communication between the client and server is not possible during the TLS/SSL handshake. After the message is modified, the Message Authentication Codes (MACs) computed by the client and server no longer match. For all subsequent handshake messages (and for the encrypted session once established), the managed device acts as a man-in-the-middle. It creates two TLS/SSL sessions, one between client and managed device, one between managed device and server. Each session contains distinct cryptographic session details.

Note The cipher suites that the system can decrypt are frequently updated and do not correspond directly to the cipher suites you can use in decryption rule conditions. For the current list of decryptable cipher suites, contact Cisco TAC .

How ServerHello and server certificate messages work

Summary

The key components involved in ServerHello and server certificate message handling are:

- Managed device: Processes ServerHello and server certificate messages and performs rule matching
- decryption rules: Define actions to take when messages match specific criteria
- Server certificate data: Certificate information used for matching against internal certificate objects
- TLS/SSL handshake: The secure connection establishment process
- Certificate authority (CA): Used for re-signing certificates in decrypt-resign actions

Workflow

These stages describe how ServerHello and server certificate message handling works:

1. The managed device receives ServerHello and server certificate messages and checks for rule matches.

   If the messages do not match any decryption rules, the managed device performs the Rule-based decryption policy default actions, on page 2179.

2. When messages match a rule that belongs to a decryption policy associated with an access control policy, the managed device performs actions based on the rule configuration.

   - **Monitor**: The TLS/SSL handshake continues to completion. The managed device tracks and logs traffic but does not decrypt it.
   - **Block or Block with Reset**: The managed device blocks the TLS/SSL session and, if configured, resets the TCP connection.
   - **Do Not Decrypt**: The TLS/SSL handshake continues to completion. The managed device does not decrypt the application data exchanged during the TLS/SSL session.

3. For decrypt actions with known key or resign methods, the managed device performs certificate matching and processing.

   - **Decrypt - Known Key**: The managed device attempts to match the server certificate data to an Internal Certificate object previously imported into the Security Cloud Control. Because you cannot generate an Internal Certificate object, and you must possess its private key, we assume you own the server on which you're using known key decryption.
   - **Decrypt - Resign**: The managed device processes the server certificate message and re-signs the server certificate with the previously imported or generated certificate authority (CA). The TLS/SSL handshake continues to completion. The managed device then uses the uploaded private key to decrypt and re-encrypt the application data exchanged during the TLS/SSL session.

   If the certificate matches a known certificate, the TLS/SSL handshake continues to completion. The managed device uses the uploaded private key to decrypt and re-encrypt the application data exchanged during the TLS/SSL session.

   If the server changes its certificate between the initial connection with the client and subsequent connections, you must import the new server certificate in the Security Cloud Control for future connections to be decrypted.



Note The system does not support mutual authentication; that is, you cannot upload a client certificate to the Cloud-Delivered Firewall Management Center and use it for **Decrypt - Resign**, **Decrypt - Replace Cert**, or **Decrypt - Known Key** decryption rule actions.

Decryption rule and policy basics

This section discusses information you should keep in mind when creating your Decryption rules and policies.

Key information for decryption policy creation

Consider this information when creating your decryption policies and rules.

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

For more information about SSL and TLS protocols, refer to a resource such as SSL vs. TLS - What's the Difference?

Guidelines for traffic decryption

Traffic that is encrypted when it passes through the system can be allowed or blocked, but it *cannot* be subjected to deep inspection or the full range of policy enforcement (such as intrusion prevention).

All encrypted connections are sent through the decryption policy to determine if they should be decrypted or blocked.

You can also configure decryption rules to block encrypted traffic of types you know you do not want on your network, such as traffic that uses the nonsecure SSL protocol or traffic with an expired or invalid certificate.

If unblocked, whether or not decrypted, traffic goes through the access control policy for a final allow or block decision.

Only *decrypted* traffic takes advantage of the system's threat defense and policy enforcement features, including:

- Advanced Malware Protection
- Security intelligence
- Threat Intelligence Director
- Application detectors
- URL and category filtering

Keep in mind that decrypting and then re-encrypting traffic adds a processing load on the device, which can reduce overall system performance.

Encrypted traffic can be allowed or blocked by policy; encrypted traffic *cannot* be inspected.

Decrypted traffic is subject to threat defense and policy enforcement; decrypted traffic can be allowed or blocked by policy.

Detects encrypted application traffic using StartTLS (SMTPS, POPS, FTPS, TelnetS, IMAPS). Also identifies certain encrypted applications based on Server Name Indication in the TLS ClientHello message or subject distinguished name from the server certificate.

Best practice for deciding when to decrypt traffic

When not to decrypt traffic

You should not decrypt traffic if doing so is forbidden by:

- Law; for example, some jurisdictions forbid decrypting financial information
- Company policy; for example, your company might forbid decrypting privileged communications
- Privacy regulations
- Traffic that uses certificate pinning (also referred to as *TLS/SSL pinning*) must remain encrypted to prevent breaking the connection

(Snort 3.) Decryption policy is *not* bypassed for any connections that match access control rules with actions of Trust, Block, or Block with Reset, unless the traffic is prefiltered. The encrypted traffic is first evaluated by decryption policy and then proceeds to the access control policy, where a final allow or block decision is made.

Encrypted traffic can be allowed or blocked on any decryption rule condition, including, but not limited to:

- Certificate status (for example, expired or invalid certificate)
- Protocol (for example, the nonsecure SSL protocol)
- Network (security zone, IP address, VLAN tag, and so on)
- URL category and reputation
- Port
- User group

Decryption rules provide a **Do Not Decrypt** action for traffic you do not want to decrypt; for more information, see Decryption rule Do Not Decrypt action, on page 2221.

Note The related information links at the end of this topic explain how some aspects of rule evaluation work. Conditions such as URL and application filtering have limitations with respect to encrypted traffic. Make sure you understand those limitations.

For more information about using URL filtering in **Do Not Decrypt** rules, see Decryption rule Do Not Decrypt action, on page 2221.

When to decrypt traffic

All encrypted traffic must be decrypted to take advantage of the system's threat protection and policy enforcement features. If your managed device has sufficient memory and processing power to decrypt traffic, you should decrypt traffic that is not prohibited by law or regulation.

If you must decide what traffic to decrypt, base your decision on the risk of allowing the traffic on your network. The system provides a flexible framework for classifying traffic using rule conditions, which include URL reputation, cipher suite, protocol, and many other factors.

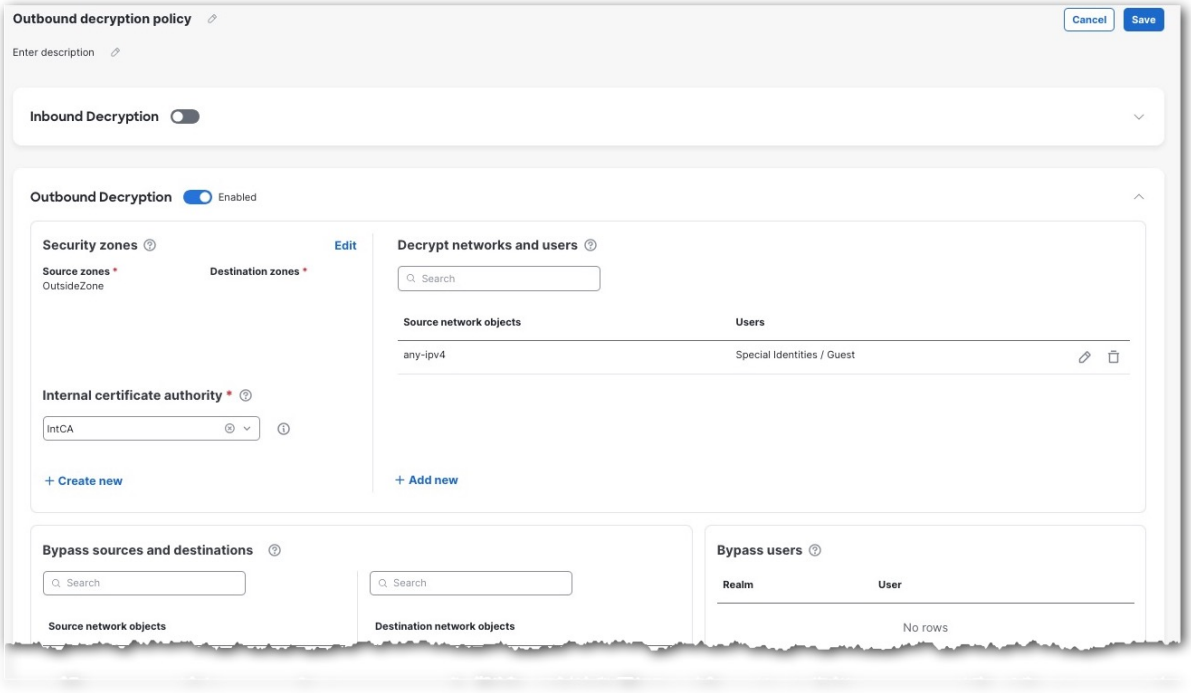
Decryption policy types

Decryption policy types include standard decryption policies and rule-based decryption policies.

Standard decryption policies

We recommend the standard decryption policy type because it's easy to set up with a wizard-like appearance, enabling you to easily pick security zones, users and networks, and other objects to use in your policy. A standard decryption policy is particularly suitable for users who are not familiar with the complexities of decryption policies.

This section provides an example of setting up a standard decryption policy.



This policy decrypts outbound traffic only. All traffic from the OutsideZone security zone on any IPv4 network is decrypted using an internal CA named `IntCA`.

Key points about decryption policies include:

- The rule is a partial example; more options are available.
- You can configure inbound criteria, outbound criteria, or both.
- In addition to objects, you can also optionally configure outbound decryption exclusions, such as:
  - Undecryptable applications (such as ones that use certificate pinning).
  - URL categories such as medical, trading, and finance.
- You can configure outbound block criteria for certificate status and TLS version.
- A standard policy has advanced policy options that are similar to rule-based policies.

Rule-based decryption policies

Rule-based decryption policies are created using a wizard that guides you through the available options for inbound decryption, outbound decryption, or both. After you create the rule-based decryption policy, you can add more rules to it, reorder rules, or make other changes to suit your needs.

A rule-based decryption policy is the most flexible but also the most potentially complicated. You can convert a standard decryption policy to a rule-based policy at any time.

Decrypt and re-sign (outgoing traffic)

Decrypt and re-sign of traffic enables the system to act as a man in the middle, intercepting, decrypting, and (if the traffic is allowed to pass) inspecting and re-encrypting traffic. It is used with outgoing traffic where the destination server is outside your protected network and requires the client to trust the firewall's certificate to complete the connection.

How decrypt and re-sign works

The Firewall Threat Defense device negotiates with the client using an internal Certificate Authority (CA) object specified in the rule and builds a TLS/SSL tunnel between the client and the Firewall Threat Defense device. At the same time, the device connects to the destination web site and creates a TLS/SSL tunnel between the server and the Firewall Threat Defense device.

Thus, the client sees the CA certificate configured for the decryption rule instead of the certificate from the destination server. The client must trust the firewall's certificate to complete the connection. The Firewall Threat Defense device then performs decryption/re-encryption in both directions for traffic between the client and the destination server.

Prerequisite

To use the **Decrypt - Resign** rule action, you must create an internal CA object using a CA file and paired private key file. You can generate a CA and private key in the system if you don't already have them.

We simplify the process of configuring the **Decrypt - Resign** rule by automatically adding **Do Not Decrypt** rules to the decryption policy based on your choices for types of traffic to exempt from encryption, such as undecryptable applications. (Typically, these applications use certificate pinning so that decrypting the connection breaks the connection.) For more information, see Create a rule-based decryption policy with outbound connection protection, on page 2165.

Note The system does not support mutual authentication; that is, you cannot upload a client certificate to the Cloud-Delivered Firewall Management Center and use it for **Decrypt - Resign**, **Decrypt - Replace Cert**, or **Decrypt - Known Key** decryption rule actions.

Incoming traffic decryption

Incoming traffic decryption protects internal servers from external attacks by intercepting and decrypting SSL/TLS encrypted traffic before it reaches the protected network.

Decryption policy requirements

This information applies only to rule-based decryption policies and rules

Note The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the Cloud-Delivered Firewall Management Center and use it for **Decrypt - Resign**, **Decrypt - Replace Cert**, or **Decrypt - Known Key** decryption rule actions.

Inbound decryption types

There are two types of inbound decryption:

- Replace Cert (default): Uses a certificate and key defined in the decryption rule to decrypt traffic. This certificate and key can be the internal server's certificate or a different certificate. You can change the certificate and key at any time. You can replace the certificate in any of the these ways:
  - API as discussed in the *Cisco Secure Firewall Threat Defense REST API Guide*.
  - Automated Certificate Management Environment (ACME).
  - **Objects** > **Object Management** > **PKI** > **Internal Certs** in the Cloud-Delivered Firewall Management Center.

  We recommend you include the certificate authority chain.
- Known Key: Use the internal server's certificate to decrypt incoming traffic. In the event the certificate changes, you must manually update it and consequently interrupt decryption until the new certificate is in place, both in the decryption policy and on the server.

Known key decryption for incoming traffic

The **Decrypt - Known Key** decryption rule action uses a server's private key to decrypt traffic. The **Decrypt - Known Key** rule action is used with incoming traffic, where the destination server is inside your protected network.

The main purpose of decrypting with a known key is to protect your servers from external attacks.

Prerequisites and requirements

To use the **Decrypt - Known Key** rule action, you must create an internal certificate object using the server's certificate file and paired private key file.

Note The system does not support mutual authentication; that is, you cannot upload a client certificate to the Cloud-Delivered Firewall Management Center and use it for **Decrypt - Resign**, **Decrypt - Replace Cert**, or **Decrypt - Known Key** decryption rule actions.

Other decryption rule actions

These topics discuss other decryption rule actions.

Decryption rule components

Each decryption rule has the following components.

State

A rule can have either of two states: enabled or disabled. By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.

Position

Rules in a decryption policy are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can match traffic by security zone, network or geographical location, VLAN, port, application, requested URL category and reputation, user, certificate, certificate subject or issuer, certificate status, cipher suite, or encryption protocol version. The use of conditions can depend on target device licenses.

Action

A rule's action determines how the system handles matching traffic. You can monitor, allow, block, or decrypt encrypted matching traffic. Decrypted traffic is subject to further inspection. Note that the system does *not* inspect blocked encrypted traffic.

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles. You can keep a record of traffic that matches a rule. You can log a connection when the system blocks an encrypted session or allows it to pass without decryption, according to the settings in an decryption policy. You can also force the system to log connections that it decrypts for further evaluation by access control rules, regardless of how the system later handles or inspects the traffic. You can log connections to the Cloud-Delivered Firewall Management Center database, as well as to the system log (syslog) or to an SNMP trap server.



Tip Properly creating and ordering decryption rules is a complex task. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the system handles traffic as you expect, the decryption policy interface has a robust warning and error feedback system for rules.

Categories

For information about using decryption rule categories (such as Applications, Category, and Cert Status), see Rule-based decryption rule conditions, on page 2201.

Rule-based decryption rule order evaluation

Rule-based decryption rule order evaluation determines the sequence in which the system evaluates decryption rules in a policy.

Rule evaluation process

This information applies only to rule-based decryption policies and rules.

When you create a rule-based decryption rule in a rule-based decryption policy , you specify its position using the **Insert** list in the rule editor. Decryption rules in an a decryption policy are numbered, starting at 1. The system matches traffic to decryption rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* decryption rule where *all* the rule's conditions match the traffic. Except in the case of Monitor rules (which log traffic but do not affect traffic flow), the system does *not* continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, VLAN, port, application, requested URL category and reputation, user, certificate, certificate distinguished name, certificate status, cipher suite, or encryption protocol version.

Each rule also has an *action*, which determines whether you monitor, block, or inspect matching encrypted or decrypted traffic with access control. Note that the system does *not* further inspect encrypted traffic it blocks. It does subject encrypted and undecryptable traffic to access control. However, access control rule conditions require unencrypted traffic, so encrypted traffic matches fewer rules.

Rules that use *specific* conditions (such as network and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your rules.



Note Proper decryption rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

In addition to ordering rules by number, you can group rules by category. By default the system provides three categories: Administrator, Standard, and Root. You can add custom categories, but you cannot delete the system-provided categories or change their order.

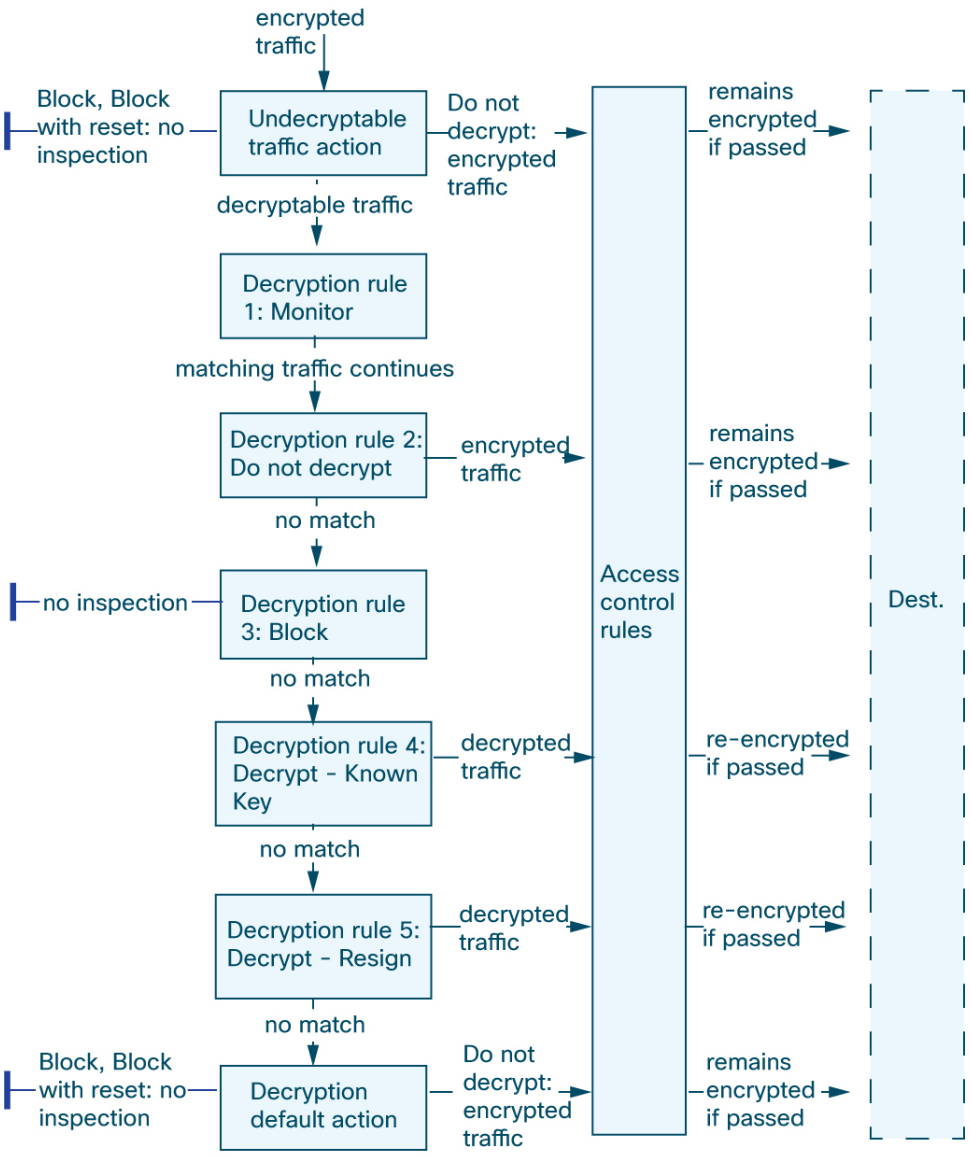
Related Topics

Best practices for access control rules, on page 2058

Multi-rule example

This reference example demonstrates how the firewall handles traffic using a decryption policy that contains multiple rules.

The following scenario summarizes the ways that decryption rules handle traffic in an inline deployment.



In this scenario, traffic is evaluated as follows:

- **Undecryptable Traffic Action** evaluates encrypted traffic first. For traffic the system cannot decrypt, the system either blocks it without further inspection or passes it for access control inspection. Encrypted traffic that does not match continues to the next rule.
- **Decryption rule 1: Monitor** evaluates encrypted traffic next. Monitor rules track and log encrypted traffic but do not affect traffic flow. The system continues to match traffic against additional rules to determine whether to permit or deny it.
- **Decryption rule 2: Do Not Decrypt** evaluates encrypted traffic third. Matching traffic is not decrypted; the system inspects this traffic with access control, but not file or intrusion inspection. Traffic that does not match continues to the next rule.

- **Decryption rule 3: Block** evaluates encrypted traffic fourth. Matching traffic is blocked without further inspection. Traffic that does not match continues to the next rule.
- **Decryption rule 4: Decrypt - Known Key** evaluates encrypted traffic fifth. Matching traffic incoming to your network is decrypted using a private key you upload. The decrypted traffic is then evaluated against access control rules. Access control rules handle decrypted and unencrypted traffic identically. The system can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the decryption rule continues to the next rule.
- **Decryption rule 5: Decrypt - Resign** is the final rule. If traffic matches, the system re-signs the server certificate with an uploaded CA certificate and decrypts traffic. The decrypted traffic is then evaluated against access control rules, which treat decrypted and unencrypted traffic the same. This additional inspection allows the system to block traffic. All remaining traffic is re-encrypted before being sent to its destination. Traffic that does not match the decryption rule rule continues to the next rule.
- **Decryption policy Default Action** handles all traffic that does not match any of the decryption rules. The default action either blocks encrypted traffic without further inspection or does not decrypt it, passing it for access control inspection.

TLS crypto acceleration

TLS crypto acceleration accelerates the following:

- TLS/SSL encryption and decryption
- VPN, including TLS/SSL and IPsec

Supported hardware and acceleration details

The following hardware models support TLS crypto acceleration:

- Secure Firewall 3100
- Secure Firewall 4200
-
- Firepower 4100/9300

  For information about TLS crypto acceleration support on Firepower 4100/9300 Firewall Threat Defense container instances, see the *FXOS Configuration Guide*.

TLS crypto acceleration is *not* supported on any virtual appliances or on any hardware except for the preceding.

Note For more information about TLS crypto acceleration and the 4100/9300, see the *FXOS Configuration Guide*.

Features *not* supported by TLS crypto acceleration include:

- Managed devices where Firewall Threat Defense container instance is enabled.

- If the inspection engine is configured to preserve connections and the inspection engine fails unexpectedly, TLS/SSL traffic is dropped until the engine restarts.

  This behavior is controlled by the **configure snort preserve-connection {enable | disable}** command.

TLS crypto acceleration guidelines and limitations

HTTP-only performance

Using TLS crypto acceleration on a managed device that is not decrypting traffic can affect performance.

Federal information processing standards (FIPS)

If TLS crypto acceleration and Federal Information Processing Standards (FIPS) are both enabled, connections with the following options fail:

- RSA keys less than 2048 bytes in size
- Rivest cipher 4 (RC4)
- Single Data Encryption Standard (single DES)
- Merkle–Damgard 5 (MD5)
- SSL v3

FIPS is enabled when you configure the Cloud-Delivered Firewall Management Center and managed devices to operate in a security certifications compliance mode. To allow connections when operating in those modes, you can configure web browsers to accept more secure options.

For more information:

- Ciphers supported by FIPS: Configure SSL settings, on page 1279.
- Security certifications compliance modes, on page 315.
- Common Criteria.

TLS heartbeat

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by RFC6520. TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

When a managed device with TLS crypto acceleration enabled encounters a packet that uses the TLS heartbeat extension, the managed device takes the action specified by the setting for **Decryption Errors** in the decryption policy's **Undecryptable Actions**:

- Block
- Block with reset

For more information, see Default handling options for undecryptable traffic, on page 2180.

To determine whether applications are using TLS heartbeat, see Troubleshoot TLS heartbeat, on page 2230.

You can configure a **Max Heartbeat Length** in a Network Analysis Policy (NAP) to determine how to handle TLS heartbeats. For more information, refer to *The SSL Preprocessor* section in the "Custom Snort 2 Intrusion Policies for Access Control" document.

TLS/SSL oversubscription

TLS/SSL oversubscription is a state where a managed device is overloaded with TLS/SSL traffic. Any managed device can experience TLS/SSL oversubscription but only managed devices that support TLS crypto acceleration provide a configurable way to handle it.

When a managed device with TLS crypto acceleration enabled is oversubscribed, any packet received by the managed device is acted on according to the setting for **Handshake Errors** in the decryption policy's **Undecryptable Actions**:

- Inherit default action
- Do not decrypt
- Block
- Block with reset

If the setting for **Handshake Errors** in the decryption policy's **Undecryptable Actions** is **Do Not decrypt** and the associated access control policy is configured to inspect the traffic, inspection occurs; decryption does *not* occur.

If a significant amount of oversubscription is occurring, you have these options:

- Upgrade your managed devices to increase TLS/SSL processing capacity.
- Change your decryption policies to add **Do Not Decrypt** rules for traffic that is not a high priority to decrypt.

Keep these guidelines in mind if your managed device has TLS crypto acceleration enabled.

View the status of TLS crypto acceleration

This task allows you to determine if TLS crypto acceleration is enabled on your managed devices.

Perform this task in the Cloud-Delivered Firewall Management Center to check the TLS crypto acceleration status for your managed devices.

Procedure

Step 1 Log in to the Cloud-Delivered Firewall Management Center.

Step 2 Click **Devices** > **Device Management**.

Step 3 Click **Edit** ( ) to edit a managed device.

Step 4 Click **Device** page.

TLS crypto acceleration status is displayed in the General section.

History for decryption policy

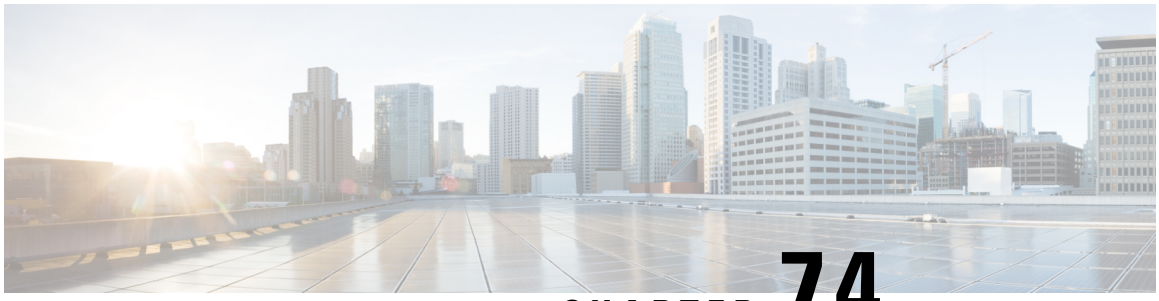
Provides a chronological overview of decryption policy features, including version support and configuration changes.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
New decryption policy user interface	December 15, 2025	Any	You have the choice of the following types of decryption policies: • Decryption policy: Also referred to as a standard decryption policy: We recommend the standard decryption policy type because it's easy to set up with a wizard-like appearance, enabling you to easily pick security zones, users and networks, and other objects to use in your policy. A standard decryption policy is particularly suited for anyone who is not proficient at understanding the ins and outs of decryption policies. • Rule-based decryption policy: Decryption policy you create using a wizard that steps you through the available options for inbound decryption, outbound decryption, or both. After you create the rule-based decryption policy, you can add more rules to it, reorder rules, or make other changes to suit your needs. A rule-based decryption policy is the most flexible but also the most potentially complicated. You can convert a standard decryption policy to a rule-based policy at any time. New/modified screens: **Policies** > **Decryption** > **Create New** > **Decryption Policy** **Policies** > **Decryption** > **Create New** > **Decryption Policy (Rule-Based)**
QUIC decryption.	20241030	7.6.0 with Snort 3	You can configure the decryption policy to apply to sessions running on the QUIC protocol. QUIC decryption is disabled by default. You can selectively enable QUIC decryption per decryption policy and write decryption rules to apply to QUIC traffic. By decrypting QUIC connections, the system can then inspect the connections for intrusion, malware, or other issues. You can also apply granular control and filtering of decrypted QUIC connections based on specific criteria in the access control policy. We modified the decryption policy Advanced Settings to include the option to enable QUIC decryption. See: Decryption Policy Advanced Options

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Easily bypass decryption for sensitive and undecryptable traffic.	20240808	7.6.0	It is now easier to bypass decryption for sensitive and undecryptable traffic, which protects users and improves performance. New decryption policies now include predefined rules that, if enabled, can automatically bypass decryption for sensitive URL categories (such as finance or medical), undecryptable distinguished names, and undecryptable applications. Distinguished names and applications are undecryptable typically because they use TLS/SSL certificate pinning, which is itself not decryptable. For outbound decryption, you enable/disable these rules as part of creating the policy. For inbound decryption, the rules are disabled by default. After the policy is created, you can edit, reorder, or delete the rules entirely. New/modified screens: **Policies** > **Access Control** > **Decryption** > **Create Decryption Policy** See: Decryption Policies
Decryption policy.	20221213	7.3.0	Feature renamed to *decryption policy* to better reflect what it does. We now enable you to configure a decryption policy with one or more **Decrypt - Resign** or **Decrypt - Known Key** rules at the same time. New/modified screens: • **Policies** > **Access Control heading** > **Decryption** (create new decryption policy) • The Create Decryption Policy dialog box now has two tab pages: **Outbound Connections** and **Inbound Connections** . • **Policies** > **Access Control heading** > **Decryption** (edit a decryption rule) **Advanced Settings** has new options discussed in TLS 1.3 decryption best practices, on page 2184 . • **Policies** > **Access Control Heading** > **Access Control** (edit an access control policy), click the word **Decryption** to associate a decryption policy with an access control policy.
TLS 1.3 decryption.	20220609	7.2.0	You can now enable TLS 1.3 decryption in an SSL policy's advanced actions. TLS 1.3 decryption requires the managed device run Snort 3. Other options are available as well; for more information, see TLS 1.3 decryption best practices, on page 2184 . New/modified screens: **SSL Policy** > **Advanced Settings**
SSL policy advanced settings.	20220609	7.2.0	SSL policy advanced settings New/modified screens: **SSL Policy** > **Advanced Settings**

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Ability to specify handling of URLs having unknown reputation.	20220609	7.0.3	For details, see URL filtering with category and reputation, on page 2086 .
ClientHello modification for **Decrypt - Known** key rules.	20220609	7.0.3	For details, see ClientHello message handling, on page 2117 .
Ability to extract the certificate in TLS 1.3 traffic to enable traffic to match URL and application criteria in access control rules.	20220609		New/modified screens: **Policies** > **Access Control** > **(edit an access control policy)** > **Advanced** link. For details, see Rule-based decryption policy advanced options, on page 2182 .
Changes to category and reputation-based URL Filtering.	20220609	7.0.3	For details, see URL filtering with category and reputation, on page 2086 .
TLS crypto acceleration cannot be disabled.	20220609	7.0.3	TLS crypto acceleration is enabled on all supported devices. On a managed device with native interfaces, TLS crypto acceleration cannot be disabled. Removed commands: • **system support ssl-hw-accel enable** • **system support ssl-hw-accel disable** • **system support ssl-hw-status**
Support for TLS crypto acceleration on one container instance on a Firepower 4100/9300 module/security engine.	20220609	7.0.3	You can now enable TLS crypto acceleration for one container instance on a module/security engine. New/modified commands: • **config hwCrypto enable** • **show crypto accelerator status**
TLS/SSL hardware acceleration is now referred to as TLS crypto acceleration .	20220609	7.0.3	The name change reflects that TLS/SSL encryption and decryption acceleration is supported on more devices. New/modified screens: **Devices** > **Device Management** > **Edit** > **Device** > **General** > **TLS Crypto Acceleration**
TLS/SSL hardware acceleration enabled by default.	20220609	7.0.3	TLS/SSL hardware acceleration is enabled by default on all supported devices but can be disabled if desired.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Extended Master Secret extension supported (see RFC 7627).	20220609	7.0.3	The TLS Extended Master Secret extension is supported for SSL policies; specifically, policies with a rule action of **Decrypt - Resign** or **Decrypt - Known Key** .
Aggressive TLS 1.3 downgrade.	20220609	7.0.3	Using the **system support ssl-client-hello-enabled aggressive_tls13_downgrade {true|false}** CLI command, you can determine the behavior for downgrading TLS 1.3 traffic to TLS 1.2.
TLS/SSL hardware acceleration introduced.	20220609	7.0.3	Certain managed device models perform TLS/SSL encryption and decryption in hardware, improving performance. By default, the feature is enabled. Affected screen: To view the status of TLS/SSL hardware acceleration , **Devices** > **Device Management** > **Device** , General page.
Category and reputation conditions supported.	20220609	7.0.3	Access control rules or SSL rules with category/reputation conditions.
SafeSearch supported.	20220609	7.0.3	The system displays an HTTP response page for connections decrypted by the SSL policy, then blocked (or interactively blocked) either by access control rules or by the access control policy default action. SafeSearch filters objectionable content and stops people from searching adult sites.
TLS/SSL policy.	20220609	7.0.3	Feature introduced.



CHAPTER 74

Standard Decryption Policies

The following topics provide details about standard decryption policies.

- Standard decryption policies, on page 2137
- Create a new standard decryption policy, on page 2139
- Standard decryption policy advanced options, on page 2157
- Decryption policy actions, on page 2159

Standard decryption policies

A standard decryption policy is a decryption policy type that is simpler to set up than other policy types and optimizes network performance by reducing processing load.

Benefits of standard decryption policies

Traffic decryption requires significant processing and decrypting too much traffic can slow down your system. The standard decryption policy type helps minimize issues commonly experienced by customers by preventing excessive traffic decryption.

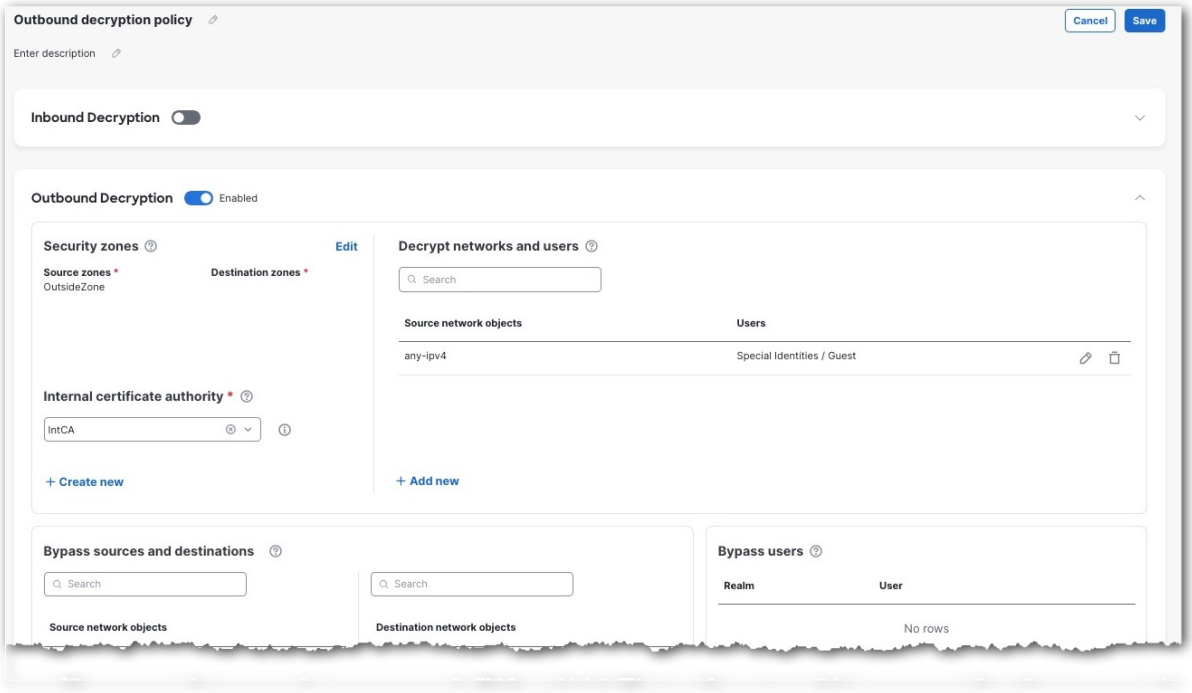
Decryption policy types

Decryption policy types include standard decryption policies and rule-based decryption policies.

Standard decryption policies

We recommend the standard decryption policy type because it's easy to set up with a wizard-like appearance, enabling you to easily pick security zones, users and networks, and other objects to use in your policy. A standard decryption policy is particularly suitable for users who are not familiar with the complexities of decryption policies.

This section provides an example of setting up a standard decryption policy.



This policy decrypts outbound traffic only. All traffic from the OutsideZone security zone on any IPv4 network is decrypted using an internal CA named `IntCA`.

Key points about decryption policies include:

- The rule is a partial example; more options are available.
- You can configure inbound criteria, outbound criteria, or both.
- In addition to objects, you can also optionally configure outbound decryption exclusions, such as:
  - Undecryptable applications (such as ones that use certificate pinning).
  - URL categories such as medical, trading, and finance.
- You can configure outbound block criteria for certificate status and TLS version.
- A standard policy has advanced policy options that are similar to rule-based policies.

Rule-based decryption policies

Rule-based decryption policies are created using a wizard that guides you through the available options for inbound decryption, outbound decryption, or both. After you create the rule-based decryption policy, you can add more rules to it, reorder rules, or make other changes to suit your needs.

A rule-based decryption policy is the most flexible but also the most potentially complicated. You can convert a standard decryption policy to a rule-based policy at any time.

Standard decryption policy deployment issues to versions earlier than 10.0.0

If you deploy a standard decryption policy to a device that runs a version earlier than 10.0.0, you can encounter compatibility issues.

- If you configured any inbound decryption rules, the rule action is changed to **Decrypt - Known Key** on those devices. For more information about the kinds of incoming decryption actions, see Incoming traffic decryption, on page 2125.
- The advanced policy option **Intelligent Decryption Bypass** requires the device to run version 7.7 or later. By default, this advanced option is disabled. If enabled, policy deployment fails when the device runs a version earlier than 7.7.

Create a new standard decryption policy

A standard decryption policy decrypts and inspects encrypted network traffic.

Decryption policy types

You can create these types of decryption policies:

- Policy to protect *inbound* traffic: Decrypt network traffic coming *into* your network. You typically do this to decrypt and inspect traffic directed to an internal server.
- Policy to protect *outbound* traffic: Decrypt network traffic going *outside* of your network to external servers.

Related Topics

Create a standard decryption policy with inbound protection, on page 2139
Create a standard decryption policy with outbound protection, on page 2144

Create a standard decryption policy with inbound protection

Create a decryption policy to decrypt and inspect traffic directed to an internal server for enhanced network security monitoring.

This task creates a standard decryption policy to decrypt network traffic coming *into* your network. You typically do this to decrypt and inspect traffic directed to an internal server. All options on this page are required.

Before you begin

Review what an inbound protection decryption policy means in Incoming traffic decryption, on page 2125.

To create a decryption policy with outbound protection instead, see Create a standard decryption policy with outbound protection, on page 2144.

Follow these steps to create a standard decryption policy with inbound protection:

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **(name of management center)**.

Step 4 In the right pane, click **Management** > **Policies**.

Step 5 Click **Create New** > **Decryption Policy**.

Step 6 In the provided fields, enter a **Name** and optional **Description**.

The following characters are not supported in decryption policy names:

- Leading period
- `#,;,{,},=,$,<,>`

Step 7 Click **Create Policy**.

Step 8 Slide **Inbound Decryption** to **Enabled** as the following figure shows.



Security zones

Security zones segment your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices. Security zones control or decrypt traffic based on its source and destination zone.

Zone configuration requirements

When configuring security zones, consider these requirements:

- If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.
- All zones used in a zone condition must be of the same type (all inline, passive, switched, or routed).
- You cannot use a zone with passive interfaces as a destination zone because devices deployed passively do not transmit traffic.
- Minimize the number of matching criteria whenever possible, especially for security zones, network objects, and port objects. If you specify multiple criteria, the system must match every combination of the criteria contents.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, the rule does not affect that device's performance.

Add security zones (inbound decryption)

Add security zones to an inbound standard decryption policy to specify Firewall Threat Defense device interfaces that send traffic to the internal server.

A security zone specifies a Firewall Threat Defense device interface that sends traffic to the internal server. Typically, for inbound protection, this will be an internal (or *DMZ*) interface.

You must choose both a source and destination security zone.

Before you begin

Complete the tasks discussed in Create a standard decryption policy with inbound protection, on page 2139.

Procedure

Step 1 Click **Edit zones** next to Security Zones.

Step 2 In the Security Zones dialog box, do any of the following:

- Select the check box next to a security zone to add to either the source or destination.
- To create a new security zone, click **Create security zone object**.
- Search for a security zone by entering text in the Search Zones field and pressing Enter.

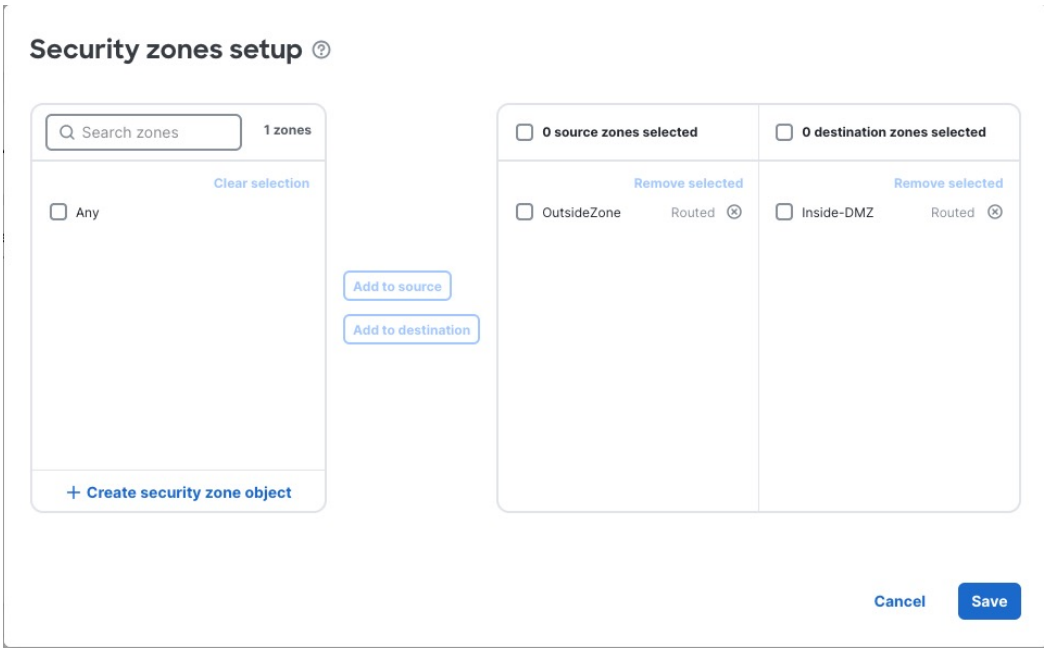
Note

Click **Help** () on any dialog box for more information.

Step 3 Click **Add to Source** to decrypt traffic that matches the source security zone or click **Add to Destination** to decrypt traffic that matches the destination security zone. If you select both source and destination zones, to be decrypted, traffic must match both zones.

Typically, your internal server should be the *destination* of an inbound decryption rule.

The following figure shows an example.



Step 4 Click **Save**.

Step 5 See Internal server details (inbound decryption), on page 2142.

Internal server details (inbound decryption)

Protect your internal servers by decrypting and optionally inspecting traffic directed to them. Add internal servers you wish to protect using network objects, and optionally, ports to specify these servers.

Network conditions and tunnel rules

Networks control or decrypt traffic based on its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Minimize the number of matching criteria, particularly for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against each combination of the specified criteria.

Behaviour of Decrypt - Replace Cert rules

When using the **Decrypt – Replace Cert** action in an inbound decryption policy, the firewall enforces strict validation of the internal server's certificate against the Trusted CA list defined within the policy. If the certificate fails validation, the firewall does not decrypt the traffic. Decryption fails if the certificate exhibits any of these conditions:

- Expired or Not-Yet-Valid: The current system time is outside the certificate's validity period.
- Untrusted: The signing Certificate Authority (CA) is not present in the Trusted CA list.
- Invalid Signature: The digital signature cannot be verified.

- Revoked or Invalid CRL: The certificate has been revoked, or the Certificate Revocation List (CRL) is invalid or unreachable.

Default handling of invalid certificates:

In a new standard decryption policy, the system generates an inbound SSL rule with the **Server Certificate Status** condition set to **Valid**.

If a server presents a certificate that does not meet these validation criteria, it fails to match the rule and is handled by the default Do Not Decrypt (DND) rule. This DND action is the default behavior for standard mode policies and is not configurable.

You can override the validation using the **Decrypt – Known** key. Refer to Override internal server validation.

Add internal servers

Add internal servers to specify certificates, networks, and ports for traffic decryption, which enables protection of internal servers and optimizes processing time.

This task discusses how to choose certificates to use to decrypt traffic coming to internal servers you want to protect. In addition to choosing a certificate, you can also specify the network and port on which the internal server is located, saving processing time.

Before you begin

Complete the tasks discussed in Create a standard decryption policy with inbound protection, on page 2139.

Follow these steps to add internal servers:

Procedure

Step 1 Click **Add new** under Internal Server Details.

Step 2 In the Internal Servers dialog box, do any of the following:

- Search for an internal certificate by entering text in the search field and pressing Enter.
- From the **Internal certificate object** list, choose an existing certificate or click **Add New** to create a new one.

  Note
  The Secure Firewall Threat Defense device servicing the traffic must trust the certificate you upload to the rule. If you upload a self-signed certificate or a certificate not trusted by a certificate authority, and you want to replace the certificate later, you must update the policy's advanced settings in any of these ways:

  - Add the certificate to the **Trusted CA certificates** list.
  - Select the **Require exact certificate match for inbound decryption** check box.

  For more information about advanced policy options, see Standard decryption policy advanced options, on page 2157.

  For more information about replacing a certificate, see the discussion of Replace Cert in Incoming traffic decryption, on page 2125.

- From the **Destination network object** list, click the network on which the internal server is located or click **Add New** to create a new one.
- (Optional.) From the **Destination port** list, click the port on which to apply the decryption rule or click **Add New** to create a new one.

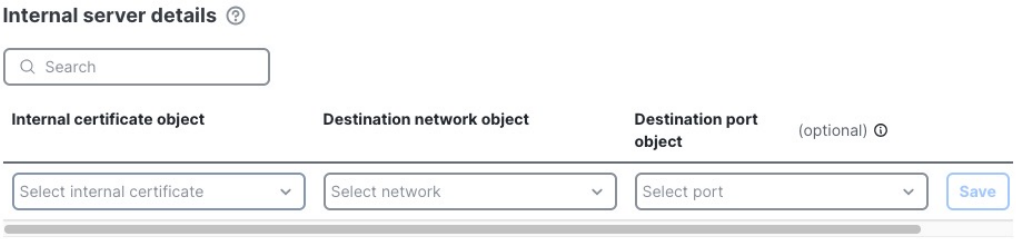
  For more information, see Upload an internal certificate for inbound protection, on page 2177.

Note

Click **Help** (?) on any dialog box for more information.

Step 3 Click **Save**.

This figure shows an example.



Step 4 Save the decryption policy by clicking **Save** at the top of the page.

Step 5 If you're finished configuring your policy, see Decryption policy actions, on page 2159.

Create a standard decryption policy with outbound protection

This task enables you to create a standard decryption policy that protects servers outside your internal network by decrypting outbound connections.

The following task discusses how to create a standard decryption policy to protect servers outside your internal network. To decrypt outbound connections, all options on this page are required except where noted. To bypass or block certain outbound connections, only security zones and bypass or block options are required.

Before you begin

Review what an outbound protection decryption policy means in Decrypt and re-sign (outgoing traffic), on page 2125.

To create a decryption policy with inbound protection instead, see Create a standard decryption policy with inbound protection, on page 2139.

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **(name of management center)**.

Step 4 In the right pane, click **Management** > **Policies**.

Step 5 Click **Create New** > **Decryption Policy**.

Step 6 In the provided fields, enter a **Name** and optional **Description**.

The following characters are not supported in decryption policy names:

- Leading period
- `#,;,{,},=,$,<,>`

Step 7 Click **Create Policy**.

Step 8 Slide **Outbound Decryption** to **Enabled** as the following figure shows.



Security zones

Security zones segment your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices. Security zones control or decrypt traffic based on its source and destination zone.

Zone configuration requirements

When configuring security zones, consider these requirements:

- If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.
- All zones used in a zone condition must be of the same type (all inline, passive, switched, or routed).
- You cannot use a zone with passive interfaces as a destination zone because devices deployed passively do not transmit traffic.
- Minimize the number of matching criteria whenever possible, especially for security zones, network objects, and port objects. If you specify multiple criteria, the system must match every combination of the criteria contents.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, the rule does not affect that device's performance.

Add security zones (outbound decryption)

Add security zones to an outbound standard decryption policy to specify Firewall Threat Defense device interfaces that send traffic to the external server.

A security zone specifies a Firewall Threat Defense device interface that sends traffic to the external server. You can configure source zones, destination zones, or both to control which traffic gets decrypted.

Before you begin

Complete the tasks discussed in Create a rule-based decryption policy with outbound connection protection, on page 2165.

Follow these steps to add security zones to an outbound decryption policy:

Procedure

Step 1 Click **Edit zones** next to Security Zones.

Step 2 In the Security Zones dialog box, do any of the following:

- Select the check box next to a security zone to add to either the source or destination.
- To create a new security zone, click **Create security zone object**.
- Search for a security zone by entering text in the Search Zones field and pressing Enter.

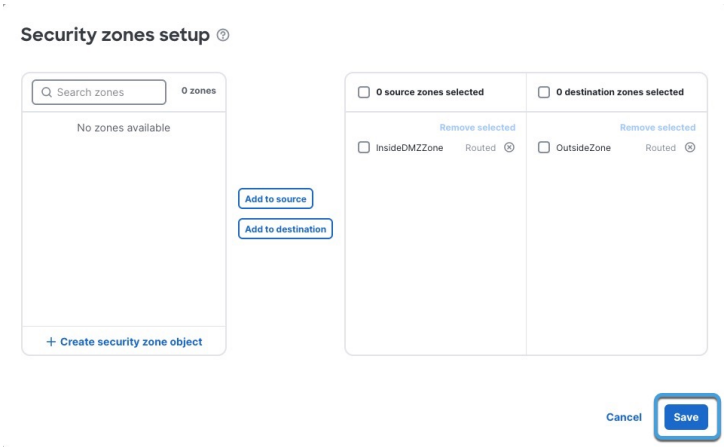
Note

Click **Help** (?) on any dialog box for more information.

Step 3 Click **Add to Source** to decrypt traffic that matches the source network or click **Add to Destination** to decrypt traffic that matches the destination network. If you select both source and destination networks, to be decrypted, traffic must match both security zones.

Typically, the server for which you're decrypting traffic should be in the *destination* zone.

The following figure shows an example.



Step 4 Click **Save**.

Step 5 If you're finished configuring your policy, see Decryption policy actions, on page 2159.

Decrypt networks and users

This option enables you to enforce decryption and deep inspection for specific network objects and from specific users and groups in Microsoft AD, LDAP, or Local realms.

Network objects

Networks control or decrypt traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions.

Users and groups

You can choose to decrypt traffic from a subset of users and groups in your identity realms and you can also choose to decrypt traffic from special identities:

- Failed Authentication: User who failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.
- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified such as users who are not downloaded by a configured realm.

For more information, see Add networks and users, on page 2147.

Add networks and users

This task allows you to enforce decryption of outbound traffic from specific networks and from users and groups in Microsoft AD, LDAP, or Local realms.

Only traffic that matches all of the networks and users you select is decrypted.

Before you begin

Complete the tasks discussed in Create a standard decryption policy with outbound protection, on page 2144.

Follow these steps to add networks and users to your decryption policy:

Procedure

Step 1 Complete the tasks discussed in Create a standard decryption policy with outbound protection, on page 2144.

Step 2 Click **Add new** next to Decrypt Networks and Users.

Step 3 From the **Source Network Objects** list, do any of the following:

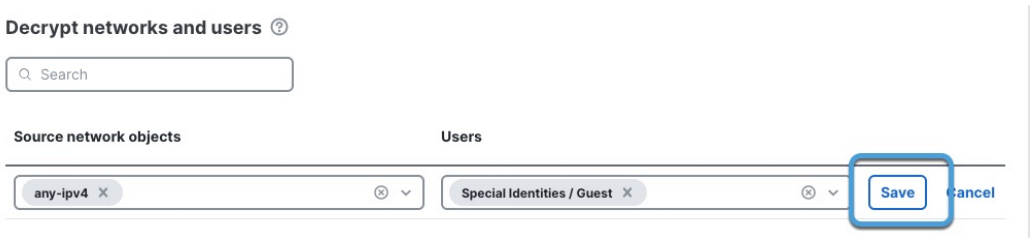
- In the provided field, enter all or part of an existing network object to filter for that object.
- Select the check box next to a network to decrypt. To decrypt traffic from all networks, click **Add New** and **Save** the row without selecting any networks.
- Click **Add new** to add another network to the list.

  For more information, see Create a network object, on page 2985.

Step 4 From the **Users** list, select the check box next to each user or group name to decrypt.

To decrypt traffic from all users, click **Add New** and **Save** the row without selecting any users.

This figure shows an example of selecting a network and users to decrypt.



Step 5 Click **Save**.

Step 6 If you're finished configuring your policy, see Decryption policy actions, on page 2159.

Internal certificate authority

An internal certificate authority resigns the traffic after it has been decrypted.

This topic discusses how to add an internal certificate authority (CA) to an outbound section of a standard decryption policy. The internal CA must be trusted by users on your network to avoid seeing untrusted certificate errors in their web browser.

For more information, see Add an internal CA for outbound protection, on page 2148.

Add an internal CA for outbound protection

This task allows you to add an internal CA to resign outgoing traffic for enhanced security in outbound protection scenarios.

This task discusses how to add an internal CA to resign outgoing traffic.

Before you begin

Complete the tasks discussed in Create a standard decryption policy with outbound protection, on page 2144.

Follow these steps to add an internal CA for outbound protection:

Procedure

In the **Internal Certificate Authority** list, do any of the following:

- Click the name of an internal CA.
- Click **Create New** to create a new internal CA.

  Click **Help** for more information.

  If you do not already have an internal CA, see:

  - Generate an internal CA for outbound protection, on page 2149
  - Upload an internal CA for outbound protection, on page 2149

Generate an internal CA for outbound protection

This task allows you to generate an internal certificate authority object to protect outbound connections when creating a decryption rule.

This task discusses how you can optionally generate an internal certificate authority when you create a decryption rule that protects outbound connections. You can also perform these tasks using **Objects** > **Object Management** as discussed in Upload a signed certificate issued in response to a CSR, on page 2990.

Before you begin

Make sure you understand the requirements for generating an internal certificate authority object as discussed in Internal certificate authority objects, on page 2987.

Procedure

Step 1 From the **Internal CA** list, click **Create New** > **Generate CA**.

Step 2 Give the internal CA a **Name** and provide a two-letter **Country Name**.

Step 3 Click **Self-Signed** or **CSR**.

For more information about these options, see Internal certificate authority objects, on page 2987.

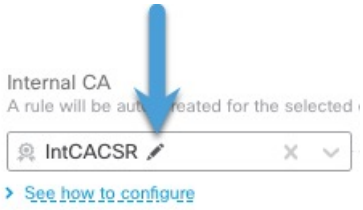
Step 4 Enter the requested information in the provided fields.

Step 5 Click **Save**.

Step 6 If you chose **CSR**, after the signing request has been completed, click **Install Certificate** as follows:

a) Repeat the preceding steps in this procedure.

b) Edit the CA from the **Internal CA** list as follows.



c) Click **Install Certificate**.

d) Follow the prompts on your screen to complete the task.

Step 7 Continue creating the policy as discussed in Create a rule-based decryption policy with outbound connection protection, on page 2165.

Upload an internal CA for outbound protection

Upload an internal certificate authority to create a decryption rule that protects outbound connections.

This task discusses how you can optionally upload an internal certificate authority when you create a decryption rule that protects outbound connections. You can also perform these tasks using **Objects** > **Object Management** as discussed in Upload a signed certificate issued in response to a CSR, on page 2990.

Before you begin

Make sure you understand the requirements for generating an internal certificate authority object as discussed in Internal certificate authority objects, on page 2987.

Follow these steps to upload an internal CA for outbound protection:

Procedure

Step 1 From the **Internal CA** list, click **Create New** > **Upload CA**.

Step 2 Give the internal CA a **Name**.

Step 3 Paste or browse to locate the certificate and its private key in the provided fields.

Step 4 If the CA has a password, select the **Encrypted** check box and enter the password in the adjacent field.

Step 5 Continue creating the policy as discussed in Create a rule-based decryption policy with outbound connection protection, on page 2165.

Bypass traffic when decrypting

You can optionally allow certain traffic to pass through the device encrypted without decryption. Bypassing certain traffic has the additional advantage that system resources are not consumed for decrypting it.

Traffic bypass methods

You can review some reasons to leave traffic encrypted here: Best practice for deciding when to decrypt traffic, on page 2122. The system provides these ways to bypass traffic when decrypting:

- Bypass source and destination networks: For example, traffic from internal servers located on an internal/DMZ network that you can trust doesn't need to be decrypted.
- Bypass users: You can bypass decryption for users and groups you trust.
- Bypass undecryptable applications: (Recommended) The typical reason to bypass outgoing traffic to applications is this traffic might use certificate pinning, which is not decryptable.

  For more information, see TLS/SSL pinning, on page 2231.
- Bypass categories: (Recommended) Bypass decrypting URL categories of sites for these reasons:
  - The categories represent applications (like personal finance or health) that might be illegal to decrypt and inspect.
  - Categories of websites Cisco has determined are low-risk.
- Intelligent decryption bypass: Bypass servers based on the threat confidence levels of clients which is determined by the Encrypted Visibility Engine (EVE) and the URL category reputation.

  All devices to which a standard decryption policy with this option enabled are deployed must run version 7.7 or later; otherwise, policy deployment fails.

Bypass sources and destinations

(Optional) Bypass traffic from source or destination network objects. For example, if traffic originates from an internal network that has trusted devices, you can opt not to decrypt that traffic. Networks control or decrypt traffic by its source and destination IP address, using inner headers.

For more information, see Add bypass sources and destinations, on page 2151.

Add bypass sources and destinations

Add bypass sources and destinations to optionally bypass outbound traffic originating from source networks or going to destination networks that you trust.

You should generally bypass traffic only if you trust the traffic from these networks.

Before you begin

Complete the tasks discussed in Create a standard decryption policy with outbound protection, on page 2144.

Follow these steps to add bypass sources and destinations:

Procedure

Step 1 Click **Add new** under Bypass sources and destinations for either source network objects or destination network objects.

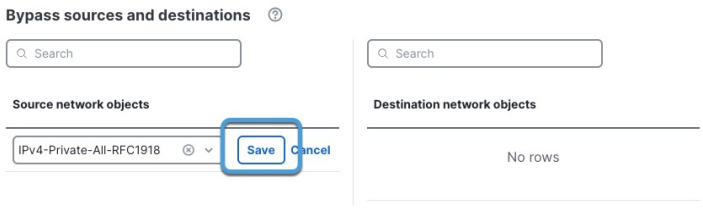
Step 2 In the dialog box that is displayed, do any of the following:

- Search for a network object by entering text in the search field and pressing Enter.
- From the **Source network object** or **Destination network object** list, choose an existing network object or click **Add New** to create a new one.

Note

Click **Help** (?) on any dialog box for more information.

The following figure shows an example.



Step 3 Click **Save**.

Step 4 If you're finished configuring your policy, see Decryption policy actions, on page 2159.

Bypass users

(Optional) You can choose to bypass decryption for users and groups in Microsoft Active Directory, LDAP, and Local realms. These users are typically ones you trust.

Bypass user configuration

You can choose to bypass a subset of users and groups in your identity realms and you can also choose to bypass from special identities:

- Failed Authentication: User who failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.
- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified such as users who are not downloaded by a configured realm.

For more information, see Add bypass users, on page 2152.

Add bypass users

Bypass decryption for users and groups in selected Microsoft AD, LDAP, and Local realms.

This task allows you to exclude specific users and groups from decryption processing by adding them to a bypass list.

Before you begin

Complete the tasks discussed in Create a standard decryption policy with outbound protection, on page 2144.

Follow these steps to add bypass users:

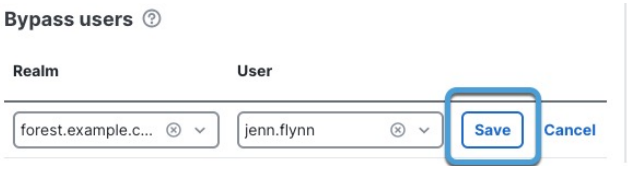
Procedure

Step 1 Click **Add new** under Bypass users.

Step 2 In the dialog box that is displayed:

- From the **Realm** list, click the name of a realm.
- From the **User** list, select the check box next to one or more users or groups for which to bypass decryption.

The following figure shows an example.



Step 3 Click **Save**.

Step 4 Repeat the preceding steps as many times as necessary to bypass more users and groups.

Bypass applications

(Optional) You can bypass decryption for applications that cannot be decrypted.

Bypass application configuration

Check the box to not decrypt traffic when re-signing the certificate is likely to cause the connection to fail.

Typically, this behavior is associated with *certificate pinning*, which is discussed in TLS/SSL certificate pinning guidelines, on page 2194.

Undecryptable applications are updated automatically in the Vulnerability Database (VDB). You can find a list of all applications on the Secure Firewall Application Detectors page; the **undecryptable** tag identifies applications Cisco determines are undecryptable.

The list of undecryptable applications is maintained by Cisco.

For more information, see Add bypass applications, on page 2153.

Add bypass applications

Bypass decryption for applications that are typically undecryptable due to their use of certificate pinning or other security mechanisms that prevent standard SSL/TLS decryption.

This task discusses how to bypass decryption for applications (typically because those applications are undecryptable due to their using certificate pinning as discussed in TLS/SSL pinning, on page 2231).

Before you begin

Complete the tasks discussed in Create a standard decryption policy with outbound protection, on page 2144.

Procedure

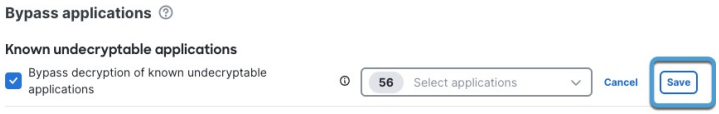
Step 1 Select the **Bypass decryption of known undecryptable applications** check box.

Step 2 Click **Edit applications**.

Step 3 From the list, do any of the following:

- Click **Select All** to select all applications that Cisco has determined are undecryptable.
- Select the check box next to the name of an application for which to bypass decryption.
- Clear the check box next to the name of an application to decrypt connections to anyway.

The following figure shows an example.



Step 4 Under Selected applications and filters, click **Add New** to add applications or application filters to bypass. Click **Help** for more information.

Step 5 Click **Save**.

Step 6 If you're finished configuring your policy, see Decryption policy actions, on page 2159.

Bypass URL categories and reputations

(Optional) Select categories and reputations, such as personal finance or health, that you do not want to decrypt.

URL categories and reputations information

URL categories and reputations are maintained by Cisco Talos.

For more information about categories, see Intelligence Categories.

For more information about reputations, see Web reputation levels.

To add URL categories and reputations to your decryption policy, see Add bypass URL categories and reputations, on page 2154.

Note Depending on the laws in your area, decryption of such traffic might be prohibited. Consult an authority in your area for more information.

Add bypass URL categories and reputations

This task allows you to bypass decryption for specific URL categories and reputations to improve performance and reduce processing overhead for trusted or sensitive content.

This task is optional and discusses how to bypass decryption of categories and reputations, based on ratings provided by Cisco Talos.

The following figure shows the list of categories and reputations we recommend.



Before you begin

Complete the tasks discussed in Create a standard decryption policy with outbound protection, on page 2144.

Follow these steps to add bypass URL categories and reputations:

Procedure

Step 1 In the Bypass URL categories section, do any of the following:

- To edit an existing category/reputation, click **Edit** ( ).
- To delete an existing category/reputation, click **Delete** ( ).
- To add a new category/reputation, click **Add new**.

- To return the list of categories and reputations to the ones Cisco recommends, click **Apply recommended settings**.

Step 2 To add or edit a category/reputation, from the first list, click the name of a category.

For more information about categories, see Intelligence Categories.

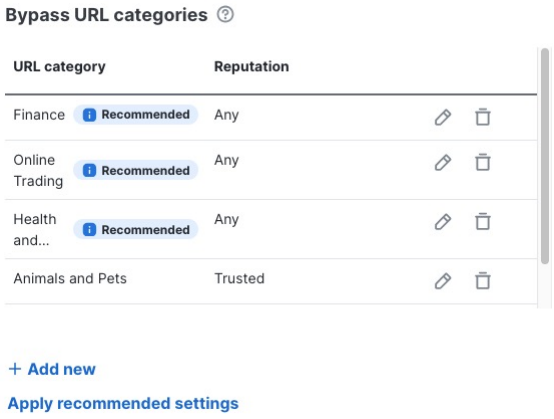
Step 3 From the second list, click the name of a reputation.

For more information about reputations, see Web reputation levels.

The following figure shows an example of using Cisco's defaults and adding the Animals and Pets category with the Trusted reputation.

Step 4 Click **Save**.

This figure shows an example.



Step 5 If you're finished configuring your policy, see Decryption policy actions, on page 2159.

Intelligent decryption bypass

Classification and setup details

To set up Intelligent Decryption Bypass for an outbound decryption rule, locate **Intelligent Decryption Bypass** and set it to **Slider enabled** ( ) to *not* decrypt traffic for very low-risk clients connecting to trusted servers, based on the threat confidence levels of clients which are determined by the Encrypted Visibility Engine (EVE) and the URL category reputation.

To view the current list of clients classified by trust level, visit Secure Firewall Application Detectors. For example, CNN.com is classified as Very Low. Make sure you understand what traffic you are choosing to *not* decrypt.

Block connections

You can choose to block any of these connection types:

- SSL and TLS versions because some are considered unsecure.

- Connections based on certificate status. For example, block outbound traffic to a server with an expired certificate because that server might not be trustworthy.

For more information, see Add block connections, on page 2156.

Add block connections

This task enables you to block outbound connections based on either old SSL or TLS versions or based on the server certificate status, enhancing security by preventing connections using insecure protocols or untrusted certificates.

This task discusses how to block connections based on either old SSL or TLS versions; or based on the server certificate status.

Note Blocking affects *all outbound connections* and takes precedence over all other bypass and decryption conditions you choose.

Before you begin

Complete the tasks discussed in Create a standard decryption policy with outbound protection, on page 2144.

Procedure

Step 1 Locate the Block connections section.

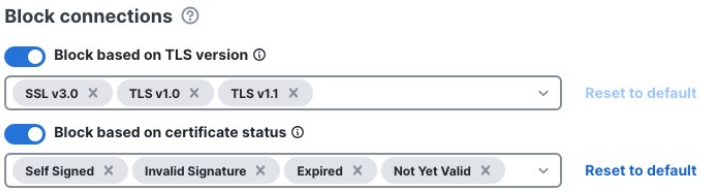
Step 2 To block outbound traffic based on its SSL or TLS protcol version, slide **Block based on TLS versions** to **Slider enabled** ( ).

Step 3 To block outbound traffic based on server certificate status, slide **Block based on certificate status** to **Slider enabled** ( ).

Step 4 You have the following options:

- From the list, select the check box next to an option to add that protocol or status to the policy.
- Click **x** next to an item to remove it from the list.
- Click **Reset to default** to return the list entries to their original values.

The following figure shows an example of blocking traffic with both TLS versions and server certificate status. Self-signed was added to the list of server certificates to block.



Step 5 If you're finished configuring your policy, see Decryption policy actions, on page 2159.

Standard decryption policy advanced options

Advanced decryption policy options allow you to fine-tune SSL/TLS decryption behavior. These optional settings control certificate matching, legacy site bypassing, protocol handling, and logging configuration. We recommend setting all advanced options to their default values unless specific requirements dictate otherwise.

Bypass legacy Cisco undecryptable sites

Enable this option to bypass decryption for websites with Distinguished Names (DNs) that Cisco has determined are undecryptable. To view the list of DNs, go to **Objects** > **Object Management** > **Distinguished Name** > **Object Groups**.

Require exact certificate match for inbound decryption

Enable this option to require the use of the internal server's certificate in the decryption policy (this option is referred to as *known key decryption*). The default is to use a different certificate, which is more convenient for replacing the policy's certificate when needed. For more information, see Incoming traffic decryption actions, on page 2223.

Override internal server validation with Decrypt – Known key

To accommodate internal servers that do not meet standard validation requirements, you can modify the policy behavior by enabling **Require exact certificate match** under **Advanced Options**.

When this option is enabled, the decryption method shifts from **Decrypt – Replace Cert** to **Decrypt – Known Key**. In this mode, the system does not perform standard CA validation. Instead, it requires the certificate presented by the internal server to be an exact match for the certificate configured in the Cloud-Delivered Firewall Management Center.

Enable adaptive TLS server identity probe

Automatically enabled when TLS 1.3 decryption is enabled. A *probe* is a partial TLS connection with the server, the purpose of which is to obtain the server certificate and cache it. (If the certificate is already cached, the probe is never established.)

If TLS 1.3 Server Identity Discovery is disabled on the access control policy with which the decryption policy is associated, we attempt to use the Server Name Indication (SNI), which is not as reliable.

The adaptive TLS server identity probe occurs on any of the following conditions as opposed to on every connection as in earlier releases:

- Certificate Issuer—Matched when the value of **Issuer DNs** in a decryption rule's DN rule condition is matched.

  For more information, refer to Conditions for distinguished name rules in decryption policies, on page 2208.

- Certificate Status—Matched when any of the **Cert Status** conditions are matched in a decryption rule.

  For more information, refer to Certificate Status Decryption rule conditions, on page 2214.

- Internal/External Certificate—Internal certificates can be matched by the certificate used in **Decrypt - Known Key** rule actions; external certificates can be matched in **Certificates** rule conditions.

For more information, refer to Known key decryption for incoming traffic, on page 2126 and Certificate rule conditions, on page 2207.

- Application ID—Can be matched by **Applications** rule conditions in either an access control policy or a decryption policy.

  For more information, refer to Application rule conditions, on page 1701.

- URL Category—Can be matched by **URLs** rule conditions in an access control policy.

  For more information, refer to Recommendation: use URL conditions for website access control , on page 1704.



Note **Enable adaptive TLS server discovery mode** is not supported on any Secure Firewall Threat Defense Virtual deployed to AWS. If you have any such managed devices managed by the Security Cloud Control, the connection event **PROBE_FLOW_DROP_BYPASS_PROXY** increments every time the device attempts to extract the server certificate.

Decryption event logging options

To expedite troubleshooting and to inform yourself about how your decryption policies are working, we recommend you enable all logging options: bypassed traffic, decrypted traffic, and blocked traffic.

Enable QUIC Decryption

Whether to apply decryption rules to connections that use the HTTP/3 over the QUIC protocol. When you decrypt QUIC connections, the system can inspect the contents of the sessions for intrusions, malware, or other issues. You can also apply granular control and filtering of decrypted QUIC connections based on specific criteria in the access control policy. QUIC support is in line with RFC 9000, 9001, 9002, 9114, 9204.

Consider the following when implementing QUIC decryption:

- On high availability or clustered devices, QUIC decryption works only if the connection remains on the same node. If the connection fails over, or is forwarded to another node, the connection drops and must be re-established. Multi-instance is supported without restrictions.
- Rules that apply to QUIC traffic would include the UDP protocol with destination port 443.
- Access control rules that apply to QUIC traffic would include the HTTP/3 or QUIC protocols, either explicitly or by implication.

The following limitations apply to QUIC decryption:

- QUIC decryption applies to Firewall Threat Defense 7.6+ only. Devices running a lower release cannot decrypt QUIC connections.
- Connections from browsers using the Chromium stack (Google Chrome, Opera, Edge) cannot be decrypted for outbound traffic. But inbound traffic from the same browsers can be decrypted.
- Connection Migration as described in RFC 9000 is not supported. The concept of Connection ID in QUIC allows endpoints to retain the same connection in the event of address change.
- Key update, session resumption, and QUIC version 2 are not supported.

- Interactive Block and Interactive Block with Reset (in access control rules) is not supported. These actions will work as Block and Block with Reset.
- The active connection-ID per connection is limited to 5. If necessary, you can modify these limits using the **system support quic-tuning** and **system support quic-tuning-reset** commands in the device CLI.

Enable TLS 1.3 decryption

Whether to apply decryption rules to TLS 1.3 connections. If you do not enable this option, the decryption rules apply to TLS 1.2 or lower traffic only. See TLS 1.3 decryption best practices, on page 2184.

Save the policy

After you have configured advanced policy options, see Decryption policy actions, on page 2159.

Add trusted CA certificates

Add trusted CA certificates to your decryption policy to ensure proper handling of your organization's certificates during traffic decryption.

This task is optional but strongly recommended. Adding issuer CAs for your organization's certificates helps avoid issues with decrypting traffic in your decryption policy.

Before you begin

Complete the tasks discussed in:

- Create a standard decryption policy with inbound protection, on page 2139
- Create a rule-based decryption policy with outbound connection protection, on page 2165

Follow these steps to add trusted CA certificates:

Procedure

Step 1 In your policy, expand **Advanced Options**.

Step 2 In the **Trusted CA certificates** list, do any of the following:

- Click the name of a trusted CA certificate from the list.
- To remove *all* certificates from the list, click .
- To return the list to its original values, click **Add recommended**.

Decryption policy actions

You can perform these actions on a decryption policy.

- Copy a decryption policy
- Convert a standard decryption policy to a rule-based decryption policy, and
- Generate a report about the decryption policy.

Related Topics

Convert a standard decryption policy to a rule-based decryption policy, on page 2160
Copy a decryption policy, on page 2160
Generate a decryption policy report, on page 2161

Convert a standard decryption policy to a rule-based decryption policy

This task allows you to use the additional customization available in a rule-based decryption policy.

After converting to a rule-based decryption policy, you cannot convert it back to a standard decryption policy. Also, you cannot convert any rule-based decryption policy to a standard decryption policy.

Before you begin

Create a decryption policy as discussed in this guide.

Follow these steps to convert a standard decryption policy to a rule-based decryption policy:

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **(name of management center)**.

Step 4 In the right pane, click **Management** > **Policies**.

Step 5 Next to the decryption policy you wish to convert, in the Actions column, click More icon (···) then click **Convert policy mode**.

Step 6 You are required to confirm the action.

The system creates a backup policy with **-backup** appended to the name of the original policy. The original policy retains its name.

Copy a decryption policy

Copy an existing decryption policy to create a new policy with the same rules and settings.

Copying a decryption policy duplicates any rules in the policy and all of the policy's settings.

Before you begin

Create a decryption policy as discussed in this guide.

Follow these steps to copy a decryption policy:

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **(name of management center)**.

Step 4 In the right pane, click **Management** > **Policies**.

Step 5 Next to the decryption policy you wish to copy, in the Actions column, click More icon (⋯) then click **Copy policy**.

Step 6 When prompted, provide a name for the policy.

Step 7 Click **Save**.

Generate a decryption policy report

This task enables you to create a comprehensive decryption policy report for analysis and documentation purposes.

This topic discusses how to create a decryption policy report that includes these items (among other things):

- List of rules
- List of cipher suites
- Policy's default action
- Trusted CA certificates

Before you begin

Create a decryption policy as discussed in this guide.

Follow these steps to generate a decryption policy report:

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

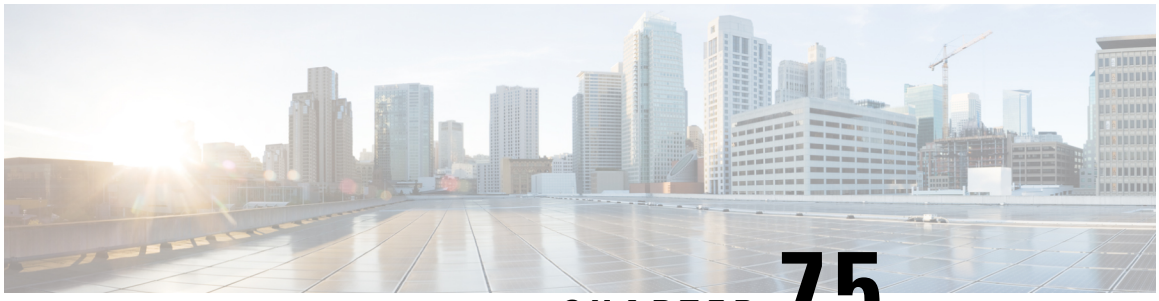
Step 2 Click **Firewall**.

Step 3 Click **Administration** > **(name of management center)**.

Step 4 In the right pane, click **Management** > **Policies**.

Step 5 Next to the decryption policy you wish to create a report, in the Actions column, click More icon (⋯) then click **Generate report**.

The system creates a PDF report and opens it for you using your default PDF application.



CHAPTER 75

Rule-Based Decryption Policies

The following topics provide details about rule-based decryption policies.

- Rule-based decryption policies, on page 2163
- Create a rule-based decryption policy , on page 2165
- Rule-based decryption policy default actions, on page 2179
- Default handling options for undecryptable traffic, on page 2180
- Rule-based decryption policy advanced options, on page 2182

Rule-based decryption policies

You can create a decryption policy using a wizard that guides you through the available options for inbound decryption, outbound decryption, or both. After you create the rule-based decryption policy, you can add more rules, reorder rules, or make other changes to suit your needs.

A rule-based decryption policy is the most flexible but also the most complex. You can convert a standard decryption policy to a rule-based policy at any time.

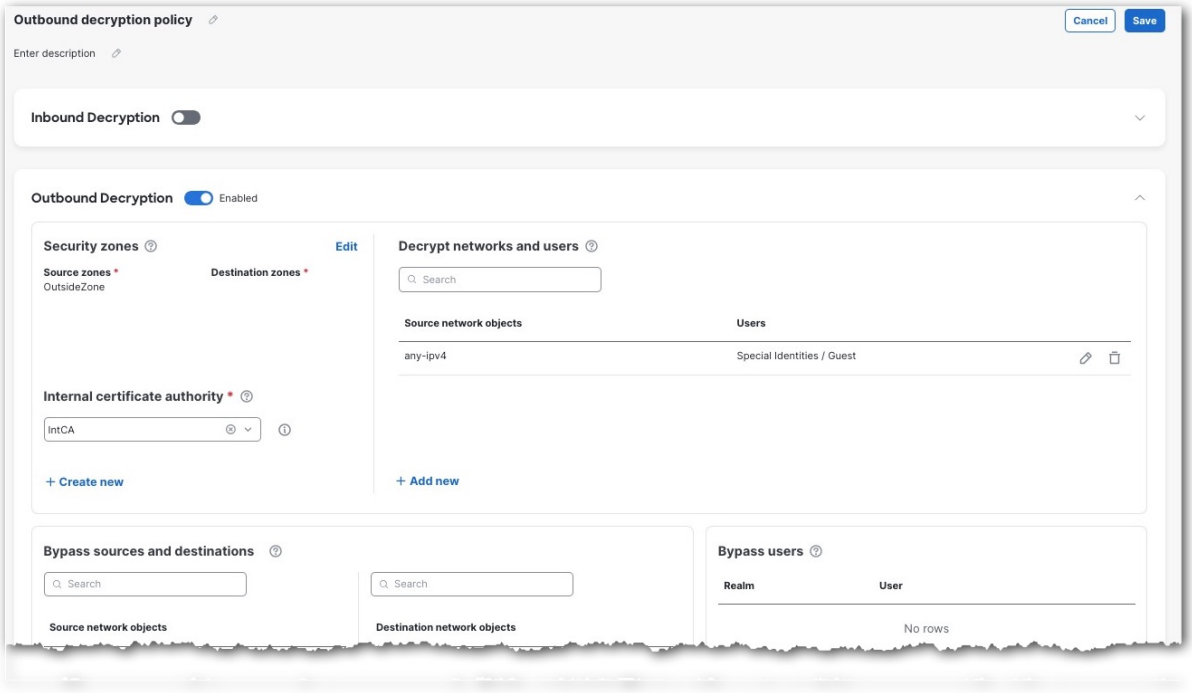
Decryption policy types

Decryption policy types include standard decryption policies and rule-based decryption policies.

Standard decryption policies

We recommend the standard decryption policy type because it's easy to set up with a wizard-like appearance, enabling you to easily pick security zones, users and networks, and other objects to use in your policy. A standard decryption policy is particularly suitable for users who are not familiar with the complexities of decryption policies.

This section provides an example of setting up a standard decryption policy.



This policy decrypts outbound traffic only. All traffic from the OutsideZone security zone on any IPv4 network is decrypted using an internal CA named `IntCA`.

Key points about decryption policies include:

- The rule is a partial example; more options are available.
- You can configure inbound criteria, outbound criteria, or both.
- In addition to objects, you can also optionally configure outbound decryption exclusions, such as:
  - Undecryptable applications (such as ones that use certificate pinning).
  - URL categories such as medical, trading, and finance.
- You can configure outbound block criteria for certificate status and TLS version.
- A standard policy has advanced policy options that are similar to rule-based policies.

Rule-based decryption policies

Rule-based decryption policies are created using a wizard that guides you through the available options for inbound decryption, outbound decryption, or both. After you create the rule-based decryption policy, you can add more rules to it, reorder rules, or make other changes to suit your needs.

A rule-based decryption policy is the most flexible but also the most potentially complicated. You can convert a standard decryption policy to a rule-based policy at any time.

Create a rule-based decryption policy

A decryption policy is a rule-based security policy that controls how encrypted network traffic is handled.

Create decryption policy to protect outbound and inbound network connections through various rule actions.

You can create any of these types of decryption policies:

- *Outbound protection* policy with rules that protect outbound connections. The destination server is outside your protected network. This type of rule has a **Decrypt - Resign** rule action. We also create additional rules with a **Do Not Decrypt** action that excludes traffic you specify (such as traffic that uses certificate pinning).

  Refer to Create a rule-based decryption policy with outbound connection protection, on page 2165 for guidance on creating a decryption rule with outbound connection.

- *Inbound protection* policy with a rule that protects inbound connections; that is, the destination server is inside your protected network. This type of rule has a **Decrypt - Known Key** rule action. We also create additional rules with a **Do Not Decrypt** action that excludes traffic you specify (such as traffic that uses certificate pinning.) These rules are disabled initially but you can modify and enable them later if you wish.

  Refer to Create a rule-based decryption policy with inbound connection protection, on page 2169 or guidance on creating a decryption rule with inbound connection.

- Other actions (including **Do Not Decrypt**, **Block**, and **Block with Reset**).

  See Create a rule-based decryption policy with other rule actions, on page 2178

Create a rule-based decryption policy with outbound connection protection

This task creates a decryption policy that protects outbound connections by intercepting and decrypting SSL/TLS traffic between internal clients and external servers, then re-encrypting the traffic with an internal CA certificate.

This task discusses how to create a decryption policy with a rule that protects outbound connections; that is, the destination server is outside your protected network. This type of rule has a **Decrypt - Resign** rule action.

When you create a decryption policy, you can create multiple rules at the same time, including multiple **Decrypt - Known Key** rules and multiple **Decrypt - Resign** rules.

If you enabled Change Management, you must create and assign a ticket before you can create a decryption policy. Before the decryption policy can be used, the ticket and all associated objects (like certificate authorities) must be approved. For more information, see Create change management tickets, on page 175 and Policies and objects that support change management, on page 169.

Before you begin

You can optionally must upload or generate an internal CA certificate for your managed device before you can create a decryption policy that protects outbound connections. You can do this in any of the following ways:

- Create an internal CA certificate object by going to **Objects** > **Object Management** > **PKI** > **Internal CAs** and referring to PKI objects, on page 2986.
- At the time you create this decryption policy.

Follow these steps to create a rule-based decryption policy with outbound connection protection:

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Administration** > **Integrations** > **Firewall Management Center** and choose **Policies** > **Access Control** > **Decryption**.

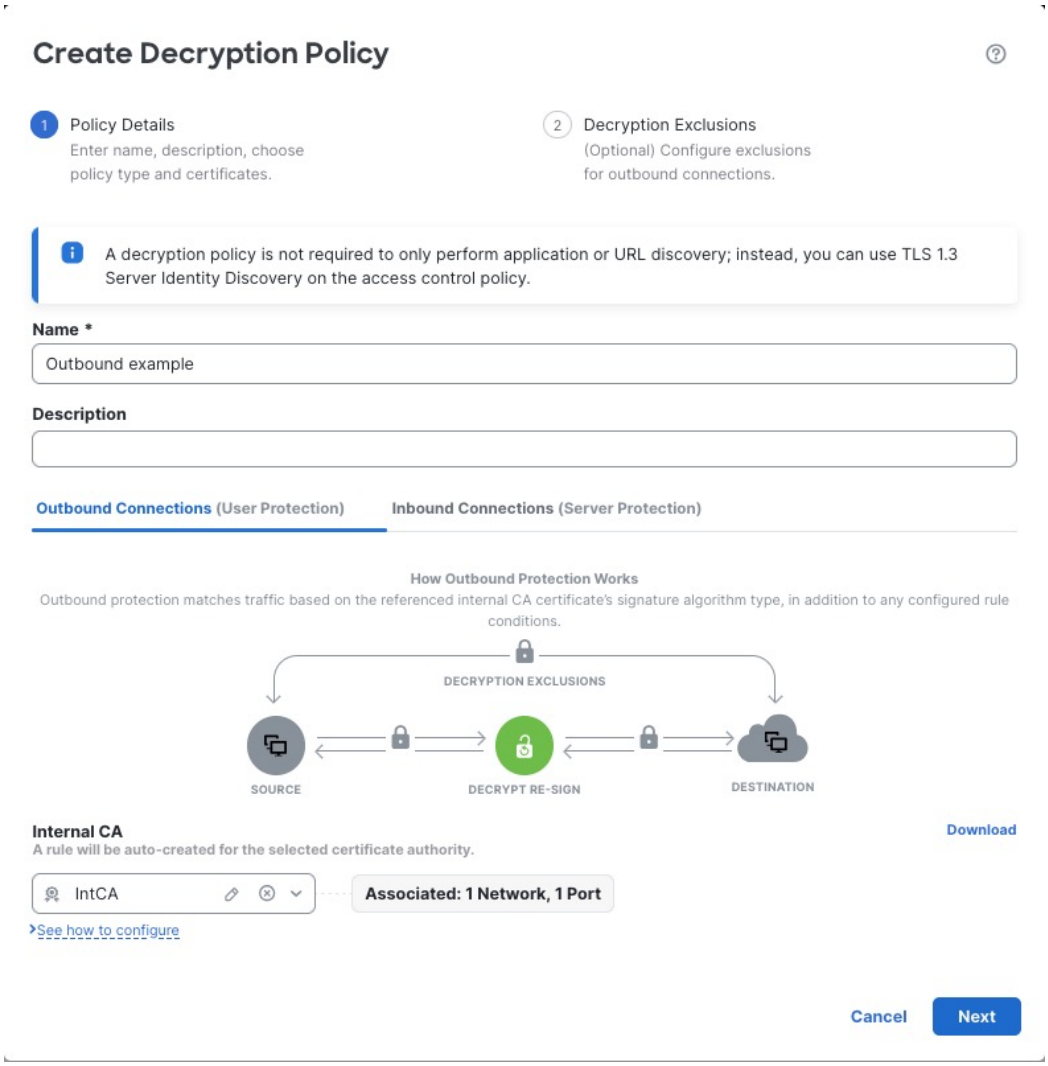
Step 3 Click **Create Decryption Policy**.

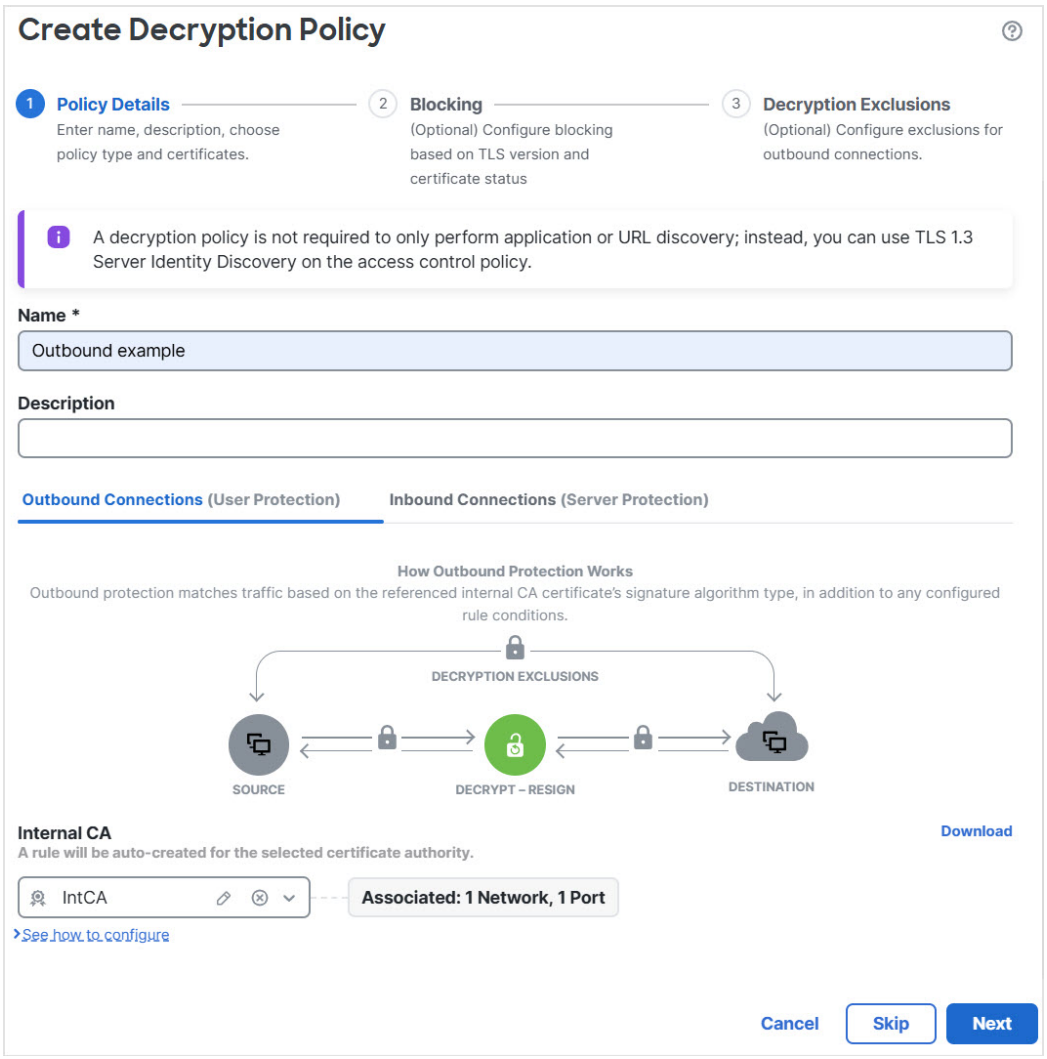
Step 4 Give the policy a unique **Name** and, optionally, a **Description**.

The following characters are not supported in decryption policy names:

- Leading period
- `#,;,{,},=,$,<,>`

Step 5 Click the **Outbound Connections** tab.


Create Decryption Policy
1 Policy Details
Enter name, description, choose policy type and certificates.
2 Decryption Exclusions
(Optional) Configure exclusions for outbound connections.
A decryption policy is not required to only perform application or URL discovery; instead, you can use TLS 1.3 Server Identity Discovery on the access control policy.
Name *
Outbound example
Description
Outbound Connections (User Protection)
Inbound Connections (Server Protection)
How Outbound Protection Works
Outbound protection matches traffic based on the referenced internal CA certificate's signature algorithm type, in addition to any configured rule conditions.
DECRYPTION EXCLUSIONS
SOURCE
DECRYPT RE-SIGN
DESTINATION
Internal CA
A rule will be auto-created for the selected certificate authority.
Download
IntCA
Associated: 1 Network, 1 Port
See how to configure
Cancel
Next



Step 6 From the **Internal CA** list, upload or choose certificates for the rules.

For more information about internal certificates, see Generate an internal CA for outbound protection, on page 2149 and Upload an internal CA for outbound protection, on page 2149.

a) (Optional.) Choose networks and ports.

For more information:

- Network rule conditions, on page 1701
- Best practice: use application-based rules over port-based rules, on page 1703

b) Click **Next**.

c) Continue with Configure Rule-based decryption policy exclusions, on page 2171.

What to do next

- Add rule conditions: Rule-based decryption rule conditions, on page 2201
- Add a default policy action: Rule-based decryption policy default actions, on page 2179
- Configure logging options for the default action .
- Set advanced policy properties: Rule-based decryption policy advanced options, on page 2182.
- Associate the decryption policy with an access control policy as described in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Deploy configuration changes.

Create a rule-based decryption policy with inbound connection protection

This task creates a decryption policy that protects inbound connections using either a **Decrypt - Known Key** rule action.

This task discusses how to create a decryption policy with a rule that protects inbound connections; that is, the destination server is inside your protected network. This type of rule has a **Decrypt - Known Key** rule action.

When you create a decryption policy, you can create multiple rules at the same time, including multiple **Decrypt - Known Key** rules and multiple **Decrypt - Resign** rules.

Before you begin

You must upload an internal certificate for your internal server before you can create a decryption policy that protects inbound connections. You can do this in any of the following ways:

- Create an internal certificate object by going to **Objects** > **Object Management** > **PKI** > **Internal Certs** and referring to PKI objects, on page 2986.
- At the time you create this decryption policy.

If you enabled Change Management, you must create and assign a ticket before you can create a decryption policy. Before the decryption policy can be used, the ticket and all associated objects (like certificate authorities) must be approved. For more information, see Create change management tickets, on page 175 and Policies and objects that support change management, on page 169.

Follow these steps to create a decryption policy with inbound connection protection:

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Tools & Services** > **Firewall Management Center** > **Policies** > **Access Control** > **Decryption**.

Step 3 Click **Create Decryption Policy**.

Step 4 Perform these steps:

a) Give the policy a unique **Name** and, optionally, a **Description**.

The following characters are not supported in decryption policy names:

- Leading period
- #,;,{,},=,$,<,>

b) From the **Internal Certificates** list, upload or choose certificates for the rules.

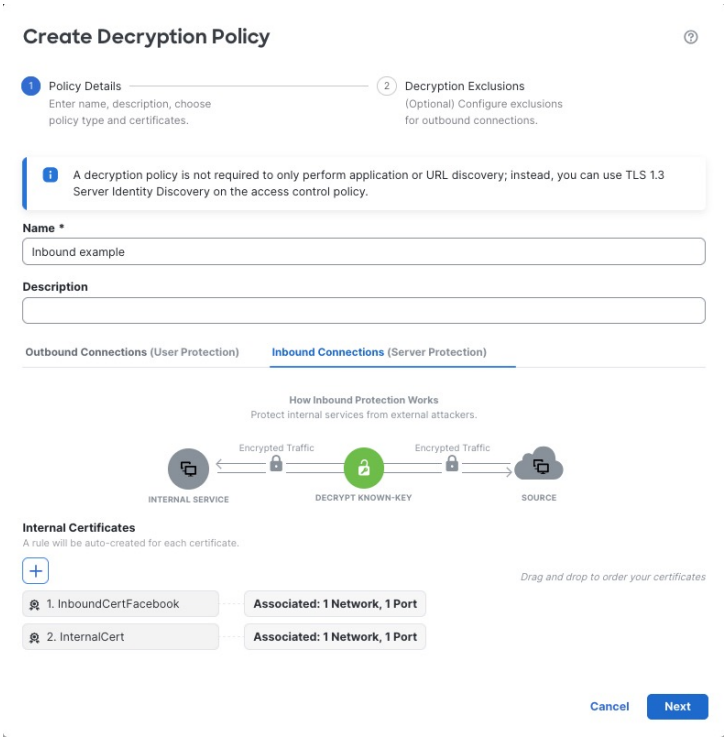
For more information about internal CA certificates, see Internal certificate authority objects, on page 2987.

c) (Optional.) Choose networks and ports.

For more information:

- Network rule conditions, on page 1701
- Best practice: use application-based rules over port-based rules, on page 1703

d) Click the **Inbound Connections** tab.



e) Click **Next**.

Step 5 Continue with Configure Rule-based decryption policy exclusions, on page 2171.

Step 6 Click **Save**.

What to do next

- Add rule conditions: Rule-based decryption rule conditions, on page 2201
- Add a default policy action: Rule-based decryption policy default actions, on page 2179

- Configure logging options for the default action.
- Set advanced policy properties: Rule-based decryption policy advanced options, on page 2182.
- Associate the decryption policy with an access control policy as described in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Deploy configuration changes.

Rule-based decryption policy block connections

Block connections to servers with unsecure TLS versions and server certificate statuses while creating a rule-based decryption policy to enhance network security.

This topic provides details about how to block connections to servers with unsecure TLS versions and server certificate statuses while creating a rule-based decryption policy . The **Block with Reset** rules are created in your decryption policy that are disabled by default.

Procedure

Step 1 Complete the tasks mentioned in:

- Create a rule-based decryption policy with outbound connection protection, on page 2165
- Create a rule-based decryption policy with inbound connection protection, on page 2169

Step 2 Configure the blocking options on the **Blocking** page.

The **Blocking** page provides blocking options. By default, all the options are *disabled* for decryption policy actions.

- **Block connections based on TLS version**—Check this check box to block connections to servers using unsecure TLS versions. By default, **SSL v3.0**, **TLS v1.0**, and **TLS v1.1** which are known to be vulnerable, are selected. You can choose other versions from the drop-down list.
- **Block connections based on server certificate status**—Check this check box to block connections to servers with unsecure server certificate statuses. By default, **Invalid Signature**, **Expired**, **Not Yet Valid**, and **Invalid Certificate** are selected. You can choose other statuses from the drop-down list.

Click **Delete** (✕) to remove the selections or click **Reset to default** to revert back to the default selections.

Step 3 Click **Next**.

What to do next

Continue with Configure Rule-based decryption policy exclusions, on page 2171.

Configure Rule-based decryption policy exclusions

This task helps you create exclusions in your decryption policy to prevent decryption of traffic that should remain encrypted for compliance, security, or functional reasons.

This task discusses how to exclude certain types of traffic from decryption. We create **Do Not Decrypt** rules in your decryption policy for this purpose. The rules are initially enabled only for an outbound decryption policy (a policy that uses the **Decrypt - Resign** policy action).

Before you begin

You must upload an internal CA certificate for your managed device before you can create a rule-based decryption policy that protects outbound connections. You can do this in any of the following ways:

- Create an internal CA certificate object by going to **Objects** > **Object Management** > **PKI** > **Internal CAs** and referring to PKI objects, on page 2986.
- At the time you create this decryption policy.

Follow these steps to configure decryption policy exclusions:

Procedure

Step 1 Complete the tasks discussed in:

- Create a rule-based decryption policy with outbound connection protection, on page 2165
- For more information, see Create a rule-based decryption policy with inbound connection protection, on page 2169

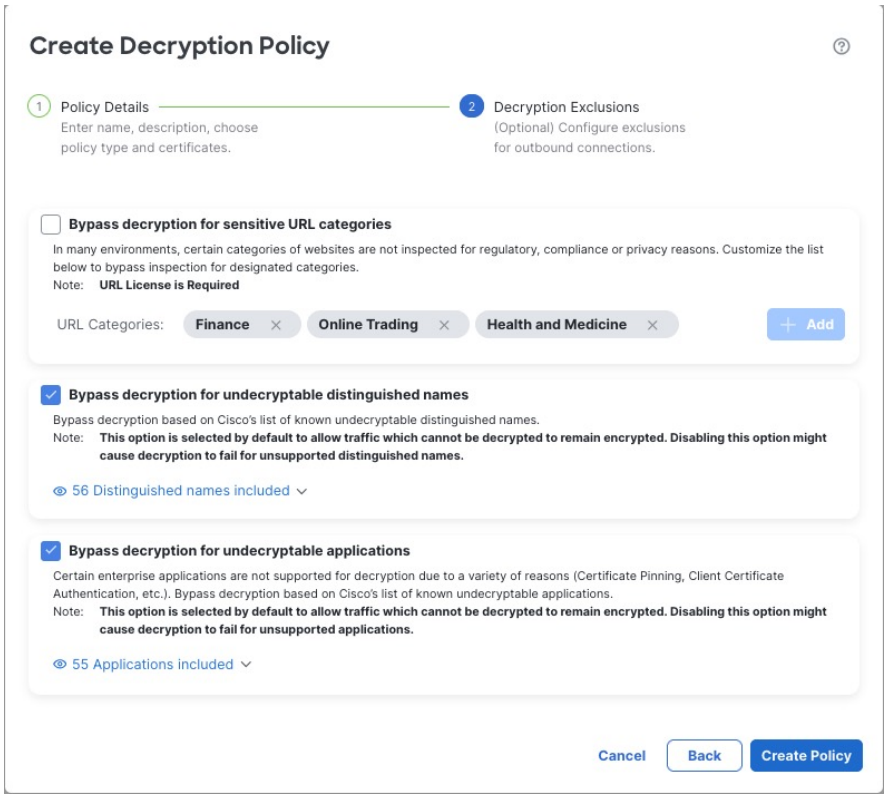
Step 2 Select the appropriate exclusion options based on your security requirements.

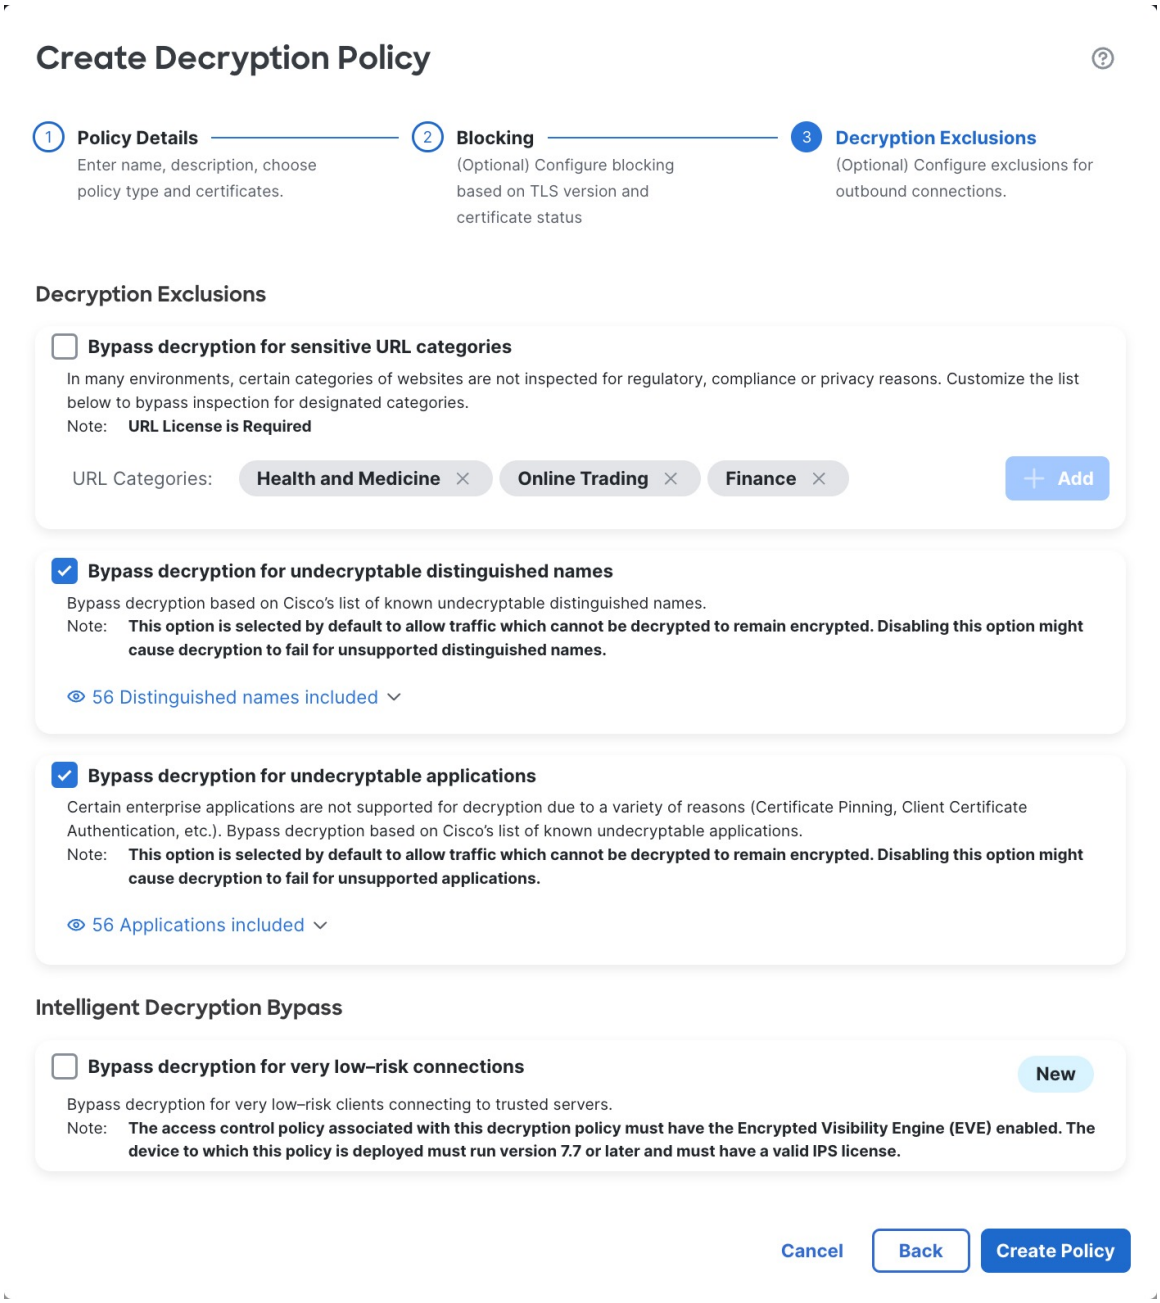
The exclusions page provides the following options. All options are *enabled* for an outbound protection policy (**Decrypt - Resign** rule action) and *disabled* for all other decryption policy actions.

Item	Description
Bypass decryption for sensitive URL categories	Check the box to not decrypt traffic from the indicated categories. Depending on the laws in your area, decryption certain traffic, such as finance or health-related, might be prohibited. Consult an authority in your area for more information. Click **Add** to add more categories. Click **Delete** (✕) to remove categories.
Bypass decryption for undecryptable distinguished names	Check the box to not decrypt traffic when re-signing the certificate is likely to cause the connection to fail. Typically, this behavior is associated with *certificate pinning*, which is discussed in TLS/SSL certificate pinning guidelines, on page 2194. The list of undecryptable distinguished names is maintained by Cisco.

Item	Description
Bypass decryption for undecryptable applications	Check the box to not decrypt traffic when re-signing the certificate is likely to cause the connection to fail. Typically, this behavior is associated with *certificate pinning*, which is discussed in TLS/SSL certificate pinning guidelines, on page 2194. Undecryptable applications are updated automatically in the Vulnerability Database (VDB). You can find a list of all applications on the Secure Firewall Application Detectors page; the **undecryptable** tag identifies applications Cisco determines are undecryptable. The list of undecryptable applications is maintained by Cisco.

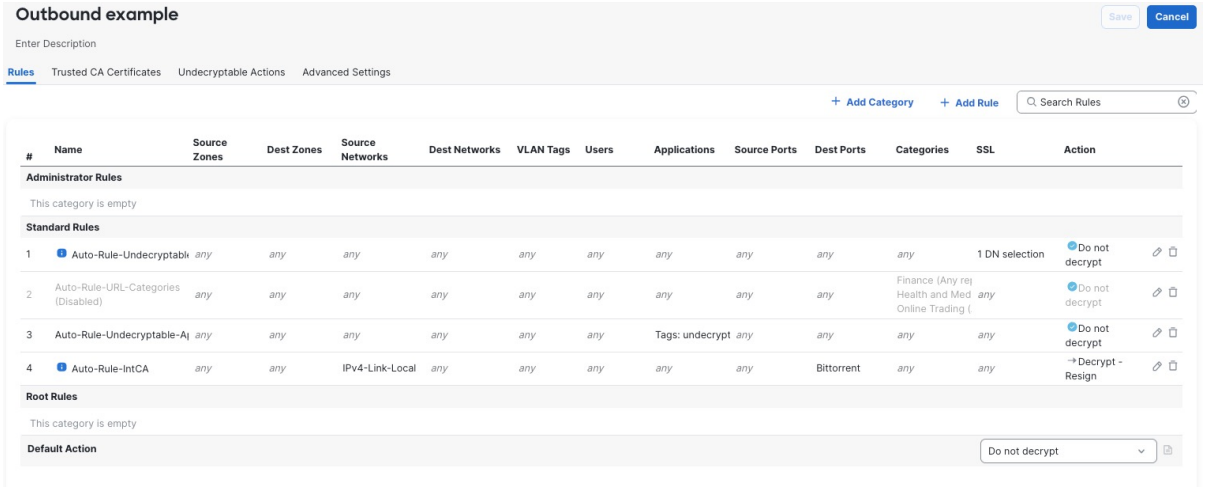
This figure shows default options.





Step 3 Click **Create Policy**.

This figure shows a sample outbound protection policy.



In the preceding example, the **Do Not Decrypt** rules corresponding to your choices for rule exclusions are automatically added before the **Decrypt - Resign** rule. The rule for sensitive URL categories is disabled because, by default, that exclusion is disabled. Had you selected the **Bypass decryption for sensitive URL categories** check box, the rule would have been enabled.

What to do next

- Add rule conditions: Rule-based decryption rule conditions, on page 2201
- Add a default policy action: Rule-based decryption policy default actions, on page 2179
- Configure logging options for the default action .
- Set advanced policy properties: Rule-based decryption policy advanced options, on page 2182.
- Associate the decryption policy with an access control policy as described in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Deploy configuration changes.

Generate an internal CA for outbound protection

This task allows you to generate an internal certificate authority object to protect outbound connections when creating a decryption rule.

This task discusses how you can optionally generate an internal certificate authority when you create a decryption rule that protects outbound connections. You can also perform these tasks using **Objects** > **Object Management** as discussed in Upload a signed certificate issued in response to a CSR, on page 2990.

Before you begin

Make sure you understand the requirements for generating an internal certificate authority object as discussed in Internal certificate authority objects, on page 2987.

Procedure

Step 1 From the **Internal CA** list, click **Create New** > **Generate CA**.

Step 2 Give the internal CA a **Name** and provide a two-letter **Country Name**.

Step 3 Click **Self-Signed** or **CSR**.

For more information about these options, see Internal certificate authority objects, on page 2987.

Step 4 Enter the requested information in the provided fields.

Step 5 Click **Save**.

Step 6 If you chose **CSR**, after the signing request has been completed, click **Install Certificate** as follows:

a) Repeat the preceding steps in this procedure.

b) Edit the CA from the **Internal CA** list as follows.



c) Click **Install Certificate**.

d) Follow the prompts on your screen to complete the task.

Step 7 Continue creating the policy as discussed in Create a rule-based decryption policy with outbound connection protection, on page 2165.

Upload an internal CA for outbound protection

Upload an internal certificate authority to create a decryption rule that protects outbound connections.

This task discusses how you can optionally upload an internal certificate authority when you create a decryption rule that protects outbound connections. You can also perform these tasks using **Objects** > **Object Management** as discussed in Upload a signed certificate issued in response to a CSR, on page 2990.

Before you begin

Make sure you understand the requirements for generating an internal certificate authority object as discussed in Internal certificate authority objects, on page 2987.

Follow these steps to upload an internal CA for outbound protection:

Procedure

Step 1 From the **Internal CA** list, click **Create New** > **Upload CA**.

Step 2 Give the internal CA a **Name**.

Step 3 Paste or browse to locate the certificate and its private key in the provided fields.

Step 4 If the CA has a password, select the **Encrypted** check box and enter the password in the adjacent field.

Step 5 Continue creating the policy as discussed in Create a rule-based decryption policy with outbound connection protection, on page 2165.

Upload an internal certificate for inbound protection

This task helps you manage certificate objects and configure decryption for secure access control by uploading an internal certificate for inbound protection.

This task discusses how to upload an internal certificate when you create a decryption rule that protects inbound connections. You can also upload the internal certificate using **Objects** > **Object Management** as discussed in Import a CA certificate and private key, on page 2988.

Before you begin

Make sure you have an internal certificate in one of the formats discussed in Internal certificate authority objects, on page 2987.

Follow these steps to upload an internal certificate for inbound protection:

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Administration** > **Integrations** > **Firewall Management Center** and choose **Policies** > **Access Control** > **Decryption**.

Step 3 Click **Create Decryption Policy**.

Step 4 Enter a name for the policy in the **Name** field and an optional description in the **Description** field.

Step 5 Click the **Inbound Connections** tab.

Step 6 From the **Internal Certificates** list, click **Add** (+).

Step 7 If an internal certificate object exists, click its name. Otherwise, click **Upload**.

a) Enter the required information.

See Add internal certificate objects, on page 2997.

Continue creating the decryption policy as discussed in Create a rule-based decryption policy with inbound connection protection, on page 2169.

Install an internal CA on client machines

To decrypt outbound traffic, the Firewall Threat Defense acts as a man-in-the-middle, first decrypting traffic (and subjecting it to deep inspection if you choose), then re-encrypting the traffic with a different internal CA. When the encrypted traffic is returned to the client, the client must trust the CA, or users see errors in their browser.

For example, the error might be:

```
www.example.com uses an invalid security certificate. The certificate is not trusted because

the issuer certificate is unknown
```

To avoid this, import the internal CA on your client machines (typically using network policies). For more information, consult a resource such as:

- Windows: Distribute Certificates to Client Computers by Using Group Policy
- macOS: Use a third party tools that is equivalent to Windows Group Policy Object (GPO). Consult the documentation provided with your system for more information.

Create a rule-based decryption policy with other rule actions

Create and manage decryption policies in the Secure Firewall Management Center to control encrypted traffic actions, such as **Do Not Decrypt**, **Block**, **Block With Reset**, or **Monitor**.

To create a decryption rule with a **Do Not Decrypt**, **Block**, **Block With Reset**, or **Monitor** rule action, create a decryption policy and edit the policy to add the rule.

When you create a rule-based decryption policy , you can create multiple rules at the same time, including multiple **Decrypt - Known Key** rules, and multiple **Decrypt - Resign** rules.

If you enabled Change Management, you must create and assign a ticket before you can create a decryption policy. Before the decryption policy can be used, the ticket and all associated objects (like certificate authorities) must be approved. For more information, see Create change management tickets, on page 175 and Policies and objects that support change management, on page 169.

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Administration** > **Integrations** > **Firewall Management Center** and choose **Policies** > **Access Control** > **Decryption**.

Step 3 Perform these actions:

a) Give the policy a unique **Name** and, optionally, a **Description**.

The following characters are not supported in decryption policy names:

- Leading period
- `#,;,{,},=,$,<,>`

b) Click **Next**.

c) The bypass page is provided for your information only; you cannot bypass traffic for other types of decryption (such as **Block**).

d) Click **Create Policy**.

e) Wait for the policy to be created.

Step 4 Click **Edit** (✎) next to the decryption policy name.

Step 5 Click **Add Rule**.

a) Give the rule a **Name**.
b) From the **Action** list, click a rule action and see one of the sections for more information:
   - Decryption rule Do Not Decrypt action, on page 2221
   - Decryption rule blocking actions, on page 2222
   - Decryption rule monitor action, on page 2220

Step 6 Click **Save**.

What to do next

- Add rule conditions: Rule-based decryption rule conditions, on page 2201
- Add a default policy action: Rule-based decryption policy default actions, on page 2179
- Configure logging options for the default action as described in *Logging Connections with a Policy Default Action* in the *Cisco Secure Firewall Management Center Administration Guide* .
- Set advanced policy properties: Rule-based decryption policy advanced options, on page 2182.
- Associate the decryption policy with an access control policy as described in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Deploy configuration changes.

Rule-based decryption policy default actions

The default action for a rule-based decryption policy determines how the system handles decryptable encrypted traffic that does not match any non-monitor rule in the policy. When you deploy a rule-based decryption policy without a decryption rules, the default action determines how all decryptable traffic on your network is handled. Note that the system does not perform any kind of inspection on encrypted traffic that the default action blocks.

To set the policy default action:

1. Log in to the Security Cloud Control if you haven't already done so.
2. Click **Administration** > **Integrations** > **Firewall Management Center** and choose **Policies** > **Access Control** > **Decryption**.
3. Click **Edit** ( ) next to the name of the decryption policy.
4. In the Default Action row, click one of the actions from the list.

Table 123: Decryption policy Default Actions

Default Action	Effect on Encrypted Traffic
Block	Block the TLS/SSL session without further inspection.

Default Action	Effect on Encrypted Traffic
Block with reset	Block the TLS/SSL session without further inspection and reset the TCP connection. Choose this option if traffic uses a connectionless protocol like UDP. In that case, the connectionless protocol tries to reestablish the connection until it is reset. This action also displays a connection reset error in the browser so the users are informed that the connection is blocked.
Do not decrypt	Inspect the encrypted traffic with access control.

Default handling options for undecryptable traffic

This reference provides information about identifying and managing undecryptable TLS/SSL traffic types in decryption policies. It covers handling handshake errors, unsupported cipher suites, compressed sessions, and session caching issues, along with user-facing actions for blocking or allowing such encrypted connections.

Table 124: Undecryptable traffic types

Type	Description	Default Action	Available Action
Compressed Session	The TLS/SSL session applies a data compression method.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
SSLv2 Session	The session is encrypted with SSL version 2. Note that traffic is decryptable if the ClientHello message is SSL 2.0, and the remainder of the transmitted traffic is SSL 3.0.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Unknown Cipher Suite	The system does not recognize the cipher suite.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Unsupported Cipher Suite	The system does not support decryption based on the detected cipher suite.	Inherit default action	Do not decrypt Block Block with reset Inherit default action

Type	Description	Default Action	Available Action
Session not cached	The TLS/SSL session has session reuse enabled, the client and server reestablished the session with the session identifier, and the system did not cache that session identifier.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Handshake Errors	An error occurred during TLS/SSL handshake negotiation.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Decryption Errors	An error occurred during traffic decryption.	Block	Block Block with Reset

When you first create a decryption policy, logging connections that are handled by the default action is disabled by default. Because the logging settings for the default action also apply to undecryptable traffic handling, logging connections handled by the undecryptable traffic actions is disabled by default.

Note that if your browser uses certificate pinning to verify a server certificate, you cannot decrypt this traffic by re-signing the server certificate. For more information, see Rule-based decryption rules guidelines and limitations, on page 2188.

Set default handling for undecryptable traffic

Set undecryptable traffic actions to handle certain types of encrypted traffic the system cannot decrypt or inspect, ensuring proper network security enforcement for encrypted connections.

You can set undecryptable traffic actions at the rule-based decryption policy level to handle certain types of encrypted traffic the system cannot decrypt or inspect. When you deploy arule-based decryption policy that contains no decryption rules, the undecryptable traffic actions determine how all undecryptable encrypted traffic on your network is handled.

Depending on the type of undecryptable traffic, you can choose to:

- Block the connection.
- Block the connection, then reset it. This option is preferrable for connectionless protocols like UDP, which keep trying to connect until the connection is blocked.
- Inspect the encrypted traffic with access control.
- Inherit the default action from the rule-based decryption policy.

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Administration** > **Integrations** > **Firewall Management Center** and choose **Policies** > **Access Control** > **Decryption**.

Step 3 Click **Edit** ( ) next to the name of the decryption policy.

Step 4 In the decryption policy editor, click **Undecryptable Actions**.

Step 5 For each field, choose either the decryption policy's default action or another action you want to take on the type of undecryptable traffic. See Default handling options for undecryptable traffic, on page 2180 and Rule-based decryption policy default actions, on page 2179 for more information.

Step 6 Click **Save** to save the policy.

What to do next

- Configure default logging for connections handled by the undecryptable traffic actions; see *Logging Connections with a Policy Default Action* in the *Cisco Secure Firewall Management Center Administration Guide*.
- Deploy configuration changes.

Rule-based decryption policy advanced options

A rule-based decryption policy 's **Advanced Settings** page contains global settings. These settings apply to all managed devices that are configured for Snort 3.

A rule-based decryption policy advanced settings are all ignored on any managed device that runs:

- A version earlier than 7.1
- Snort 2

Block flows requesting ESNI

Encrypted Server Name Indication (ESNI (link to draft proposal)) is a way for a client to tell a TLS 1.3 server what the client is requesting. Because the SNI is encrypted, the system cannot identify the server. You can choose to block these connections.

Disable HTTP/3 advertisement

This option strips HTTP/3 (RFC 9114) from the ClientHello in TCP connections. HTTP/3 is part of the QUIC transport protocol, not the TCP transport protocol. Blocking clients from advertising HTTP/3 provides protection against attacks and evasion attempts potentially burried within QUIC connections.

Propagate untrusted server certificates to clients

This setting applies only to traffic matching a **Decrypt - Resign** rule action.

Enable this option when the server certificate is untrusted. The managed device substitutes its certificate authority (CA) for the server certificate. An *untrusted* server certificate is one that is not listed as a trusted CA in the Security Cloud Control(**Objects**, **Objects** > **Object Management** > **PKI** > **Trusted CAs**).

Enable TLS 1.3 decryption

Whether to apply decryption rules to TLS 1.3 connections. If you do not enable this option, the decryption rules apply to TLS 1.2 or lower traffic only. See TLS 1.3 decryption best practices, on page 2184.

Enable adaptive TLS server identity probe

Automatically enabled when TLS 1.3 decryption is enabled. A *probe* is a partial TLS connection with the server, the purpose of which is to obtain the server certificate and cache it. (If the certificate is already cached, the probe is never established.)

If TLS 1.3 Server Identity Discovery is disabled on the access control policy with which the decryption policy is associated, we attempt to use the Server Name Indication (SNI), which is not as reliable.

The adaptive TLS server identity PROBE occurs on any of these conditions as opposed to on every connection as in earlier releases:

- Certificate Issuer—Matched when the value of **Issuer DNs** in a decryption rule's DN rule condition is matched.

  For more information, see Conditions for distinguished name rules in decryption policies, on page 2208.

- Certificate Status—Matched when any of the **Cert Status** conditions are matched in a decryption rule.

  For more information, see Certificate Status Decryption rule conditions, on page 2214.

- Internal/External Certificate—Internal certificates can be matched by the certificate used in **Decrypt - Known Key** rule actions; external certificates can be matched in **Certificates** rule conditions.

  For more information, see Known key decryption for incoming traffic, on page 2126 and Certificate rule conditions, on page 2207.

- Application ID—Can be matched by **Applications** rule conditions in either an access control policy or a decryption policy.

  For more information, see Application rule conditions, on page 1701.

- URL Category—Can be matched by **URLs** rule conditions in an access control policy.

  For more information, see Recommendation: use URL conditions for website access control , on page 1704.



Note **Enable adaptive TLS server discovery mode** is not supported on any Secure Firewall Threat Defense Virtual deployed to AWS. If you have any such managed devices managed by the Security Cloud Control, the connection event **PROBE_FLOW_DROP_BYPASS_PROXY** increments every time the device attempts to extract the server certificate.

Enable QUIC Decryption

Whether to apply decryption rules to connections that use the HTTP/3 over the QUIC protocol. When you decrypt QUIC connections, the system can inspect the contents of the sessions for intrusions, malware, or other issues. You can also apply granular control and filtering of decrypted QUIC connections based on specific criteria in the access control policy. QUIC support is in line with RFC 9000, 9001, 9002, 9114, 9204.

Consider the following when implementing QUIC decryption:

- On high availability or clustered devices, QUIC decryption works only if the connection remains on the same node. If the connection fails over, or is forwarded to another node, the connection drops and must be re-established. Multi-instance is supported without restrictions.
- Rules that apply to QUIC traffic would include the UDP protocol with destination port 443.
- Access control rules that apply to QUIC traffic would include the HTTP/3 or QUIC protocols, either explicitly or by implication.

The following limitations apply to QUIC decryption:

- QUIC decryption applies to Firewall Threat Defense 7.6+ only. Devices running a lower release cannot decrypt QUIC connections.
- Connections from browsers using the Chromium stack (Google Chrome, Opera, Edge) cannot be decrypted for outbound traffic. But inbound traffic from the same browsers can be decrypted.
- Connection Migration as described in RFC 9000 is not supported. The concept of Connection ID in QUIC allows endpoints to retain the same connection in the event of address change.
- Key update, session resumption, and QUIC version 2 are not supported.
- Interactive Block and Interactive Block with Reset (in access control rules) is not supported. These actions will work as Block and Block with Reset.
- The active connection-ID per connection is limited to 5. If necessary, you can modify these limits using the **system support quic-tuning** and **system support quic-tuning-reset** commands in the device CLI.

TLS 1.3 decryption best practices

Both the rule-based decryption policy and the access control policy have advanced options that affect how traffic is handled, whether the traffic is being decrypted or not.

The advanced options are:

- Decryption policy:
  - TLS 1.3 decryption
  - TLS adaptive server identity probe
- Access control policy: TLS 1.3 Server Identity Discovery

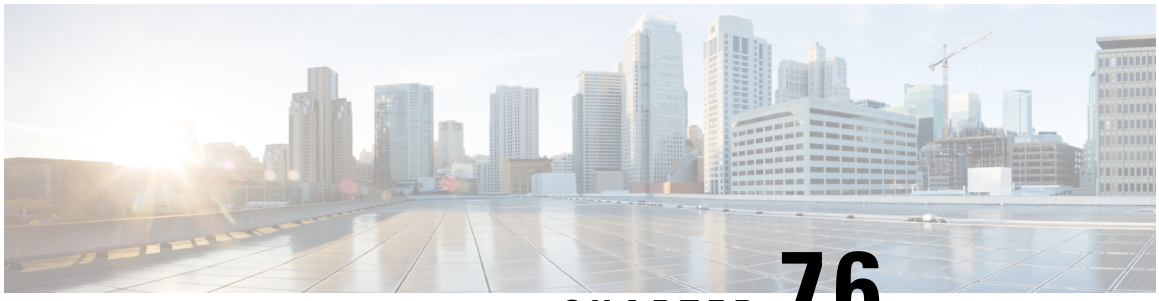
  The access control policy setting takes precedence over the decryption policy setting.

Use the following table to decide which option to enable:

TLS adaptive server identity probe setting (decryption policy)	TLS 1.3 Server Identity Discovery setting (access control policy)	Result	Recommended when
Enabled	Disabled	Adaptive probe sent if decryption policy contains *any* rule conditions specified in Rule-based decryption policy advanced options, on page 2182 *and* if the server certificate is not cached.	• You're not using application or URL conditions in access control rules • You're decrypting traffic
Enabled	Enabled	Probe is always sent if the server certificate is not cached.	Use only if your access control rules have URL or application conditions
Disabled	Enabled	Probe is always sent if the server certificate is not cached.	Not recommended.
Disabled	Disabled	Probe is never sent.	Very limited usefulness; use only if not decrypting traffic and not using application or URL conditions in the access control rule

Note A cached TLS server's certificate is available to all Snort instances on a particular Firewall Threat Defense. The cache can be cleared with a CLI command and is automatically cleared when the device is rebooted.

For more information, see the discussion of TLS server identity discovery on secure.cisco.com.



CHAPTER 76

Rule-Based Decryption Rules

The following topics discuss how to manage rule-based decryption rules.

- Rule-based decryption, on page 2187
- Requirements and prerequisites for Rule-based decryption rules, on page 2187
- Rule-based decryption rules guidelines and limitations, on page 2188
- Best practice for Decryption policy rule order, on page 2196
- Rule-based decryption rule traffic handling, on page 2197
- Rule-based decryption rule conditions, on page 2201
- Rule-based decryption rule actions, on page 2220
- Monitor TLS/SSL hardware acceleration, on page 2226
- Troubleshoot Rule-based decryption rules, on page 2228
- Troubleshoot using crypto archives, on page 2236

Rule-based decryption

Decryption rules are a granular method that

- handle encrypted traffic across multiple managed devices
- provide options to block traffic without inspection, inspect without decrypting, or decrypt for access control analysis, and
- support access control rule evaluation.

Requirements and prerequisites for Rule-based decryption rules

This reference provides the domain support and user role requirements necessary for Rule-based decryption rules.

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Rule-based decryption rules guidelines and limitations

Keep these points in mind when setting up your rule-based decryption rule. Properly configuring rule-based decryption rule is a complex task, but one that is essential to building an effective deployment that handles encrypted traffic. Many factors influence how you configure rules, including certain application behavior that you cannot control.

In addition, rules can preempt each other, require additional licenses, or contain invalid configurations. Thoughtfully configured rules can also reduce the resources required to process network traffic. Creating overly complex rules and ordering rules the wrong way can adversely affect performance.

For detailed information, see Best practices for access control rules, on page 2058.

Related Topics

Best practices for access control rules, on page 2058
Guidelines for using TLS/SSL decryption, on page 2188
Rule-based decryption rules unsupported features, on page 2190
Guidelines for TLS/SSL Do Not Decrypt configuration, on page 2190
Best practices for TLS/SSL incoming traffic decryption, on page 2191
TLS/SSL Decrypt - Resign guidelines, on page 2192
TLS/SSL block guidelines, on page 2194
TLS/SSL certificate pinning guidelines, on page 2194
TLS/SSL heartbeat guidelines, on page 2195
TLS/SSL anonymous cipher suite limitation, on page 2195
TLS/SSL normalizer guidelines, on page 2195
Other Rule-based decryption rule guidelines, on page 2195
Rule and other policy warnings
Best practice for Decryption policy rule order, on page 2196

Guidelines for using TLS/SSL decryption

General guideline

Rule-based decryption rules require processing overhead that can impact performance. Determine which traffic must be decrypted and subjected to deep inspection before you set up any policies or rules.

- You cannot decrypt traffic on a device that uses passive interface.
- You cannot decrypt traffic on a device that uses inline tap interface.
- You cannot decrypt traffic on a device that uses TCP state bypass (only if the interfaces are in routed mode).

Guidelines for undecryptable traffic

We can determine that certain traffic is not decryptable either because the website itself is not decryptable or because the website uses TLS/SSL pinning, which effectively prevents users from accessing a decrypted site without errors in their browser.

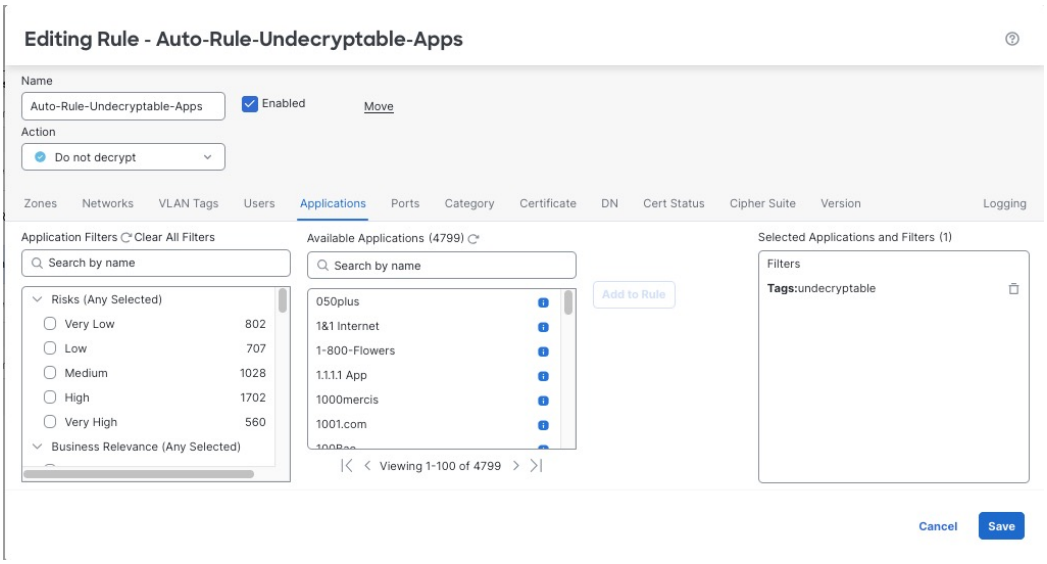
For more information about certificate pinning, see TLS/SSL pinning, on page 2231.

We maintain the list of these sites as follows:

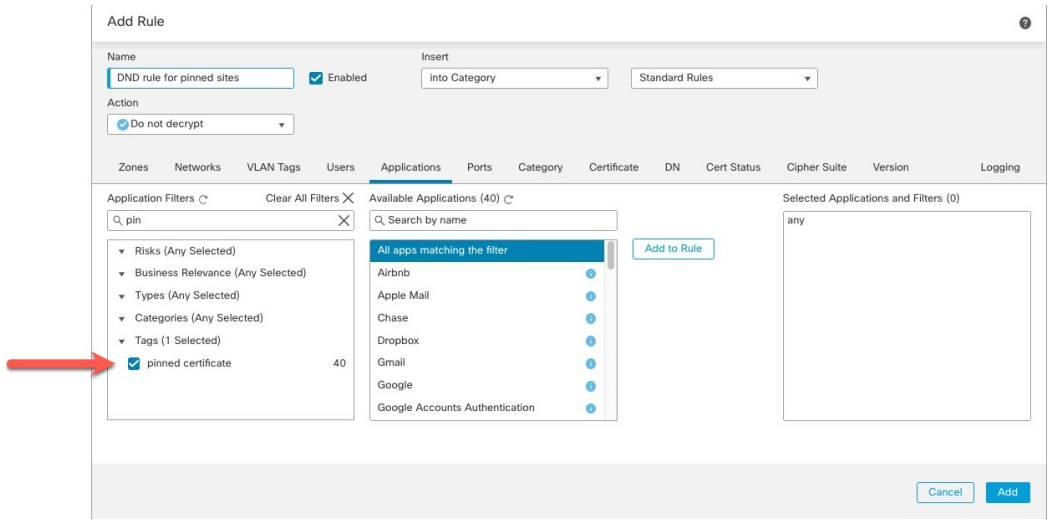
- A Distinguished Name (DN) group named **Cisco-Undecryptable-Sites**
- The **pinned certificate** or **undecryptable** application filter

If you are decrypting traffic and you do not want users to see errors in their browsers when going to these sites, we recommend you set up a **Do Not Decrypt** rule toward the bottom of your decryption rules.

If you use the decryption policy wizard to create a policy for outbound traffic protection, a **Do Not Decrypt** rule for pinned certificates is created for you as the following example shows.



An example of setting up a **pinned certificate** application filter follows.



Rule-based decryption rules unsupported features

This reference provides guidance for handling RC4, DES, and TLS 1.3 limitations in the Cisco Firepower System, including configuration requirements for decryption policies when dealing with unsupported cipher suites and interface types.

RC4 cipher not unsupported

The Rivest Cipher 4 (also referred to as *RC4* or *ARC4* cipher is known to have vulnerabilities and is considered unsecure.
decryption policies identify the RC4 cipher suite as unsupported; you should configure the **Unsupported Cipher Suite** action in the decryption policies policy's **Undecryptable Actions** page to match your organization's requirements. For more information, see Default handling options for undecryptable traffic, on page 2180.

Passive, inline tap mode, and SPAN interfaces not supported

TLS/SSL traffic cannot be decrypted on passive, inline tap mode, or SPAN interfaces.

Guidelines for TLS/SSL Do Not Decrypt configuration

You should not decrypt traffic if doing so is forbidden by:

- Law; for example, some jurisdictions forbid decrypting financial information
- Company policy; for example, your company might forbid decrypting privileged communications
- Privacy regulations
- Traffic that uses certificate pinning (also referred to as *TLS/SSL pinning*) must remain encrypted to prevent breaking the connection

Encrypted traffic can be allowed or blocked on any decryption rule condition, including, but not limited to:

- Certificate status (for example, expired or invalid certificate)
- Protocol (for example, the nonsecure SSL protocol)
- Network (security zone, IP address, VLAN tag, and so on)

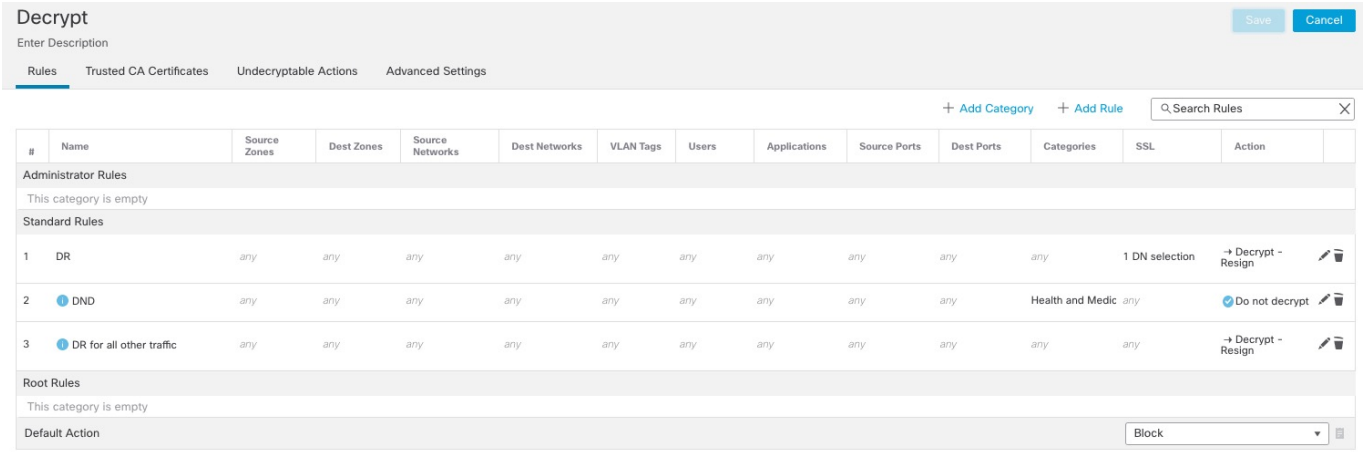
- URL category and reputation
- Port
- User group

Limitations of categories in Do Not Decrypt rules

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is not based on the declared company vertical, intent, or security. While we strive to continuously update and improve URL filtering categories, it is not an exact science. Some websites are not categorized at all and it's possible some websites might be improperly categorized.

Avoid overusing categories in do not decrypt rules to avoid decrypting traffic without a reason; for example, the Health and Medicine category includes the WebMD website, which does not threaten patient privacy.

Following is a sample decryption policy that can prevent decryption for websites in the Health and Medicine categories but allows decryption for WebMD and everything else. General information about decryption rules can be found in Guidelines for using TLS/SSL decryption, on page 2188.



Note Don't confuse URL filtering with application detection, which relies on reading some of the packet from a website to determine more specifically what it is (for example, Facebook Message or Salesforce). For more information, see Recommendations for application control, on page 2054.

Best practices for TLS/SSL incoming traffic decryption

Certificate and key association
Known key rule action

You can associate one or more server certificates and paired private keys with a **Decrypt - Known Key** rule action.

If traffic matches the rule, and the certificate used to encrypt the traffic matches the certificate associated with the action, the system uses the appropriate private key to obtain the session encryption and decryption keys. Because you must have access to the private key, this action is best suited to decrypt traffic incoming to servers your organization controls.

Restrictions and considerations

Note these restrictions when configuring decryption rules:

- Anonymous cipher suite unsupported:

  By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

  You cannot use the **Decrypt - Resign**, **Decrypt - Replace Cert** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

- Cannot match on Distinguished Name or Certificate:

  You cannot match on **Distinguished Name** or **Certificate** rule conditions when creating a decryption rule with a **Decrypt - Known Key** action. The assumption is that if this rule matches traffic, the certificate, subject DN, and issuer DN already match the certificate associated with the rule.

- Do not use Version or Cipher Suite rule conditions.



Important *Never* use either **Cipher Suite** or **Version** rule conditions in a rule with a **Decrypt - Resign** or **Decrypt - Known Key** rule action. The use of these conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

- Elliptic Curve Digital Signature Algorithm (ECDSA) certificate results in blocked traffic:

  (TLS 1.3 decryption enabled only.) If you use an ECDSA certificate with a **Decrypt - Known Key** rule action, matching traffic will be blocked. To avoid this, use a certificate with another type of certificate.

TLS/SSL Decrypt - Resign guidelines

You can associate one internal Certificate Authority (CA) certificate and private key with the **Decrypt - Resign** action. If traffic matches this rule, the system re-signs the server certificate with the CA certificate, then acts as a man-in-the-middle. It creates two TLS/SSL sessions, one between client and managed device, one between managed device and server. Each session contains different cryptographic session details, and allows the system to decrypt and reencrypt traffic.

- Use the **Decrypt - Resign** rule action for decrypting *outgoing* traffic, as opposed to incoming traffic for which we recommend the **Decrypt - Known Key** rule action.

  For more information about **Decrypt - Known Key**, see .

- Always check the **Replace Key Only** check box when you set up a **Decrypt - Resign** rule action.

  When a user browses to a web site that uses a *self-signed* certificate, the user sees a security warning in the web browser and is aware that they are communicating with an unsecure site.

  When a user browses to a web site that uses a trusted certificate, the user does not see a security warning.

Certificate algorithm compatibility

If you configure a rule with the **Decrypt - Resign** action, the rule matches traffic based on the referenced internal CA certificate's signature algorithm type, in addition to any configured rule conditions. Because you associate one CA certificate with a **Decrypt - Resign** action, you cannot create a decryption rule that decrypts multiple types of outgoing traffic encrypted with different signature algorithms. In addition, any external certificate objects and cipher suites you add to the rule must match the associated CA certificate encryption algorithm type.

- For example, outgoing traffic encrypted with an elliptic curve (EC) algorithm matches a **Decrypt - Resign** rule only if the action references an EC-based CA certificate; you must add EC-based external certificates and cipher suites to the rule to create certificate and cipher suite rule conditions.
- Similarly, a **Decrypt - Resign** rule that references an RSA-based CA certificate matches only outgoing traffic encrypted with an RSA algorithm; outgoing traffic encrypted with an EC algorithm does not match the rule, even if all other configured rule conditions match.

Guidelines and limitations

Anonymous cipher suite unsupported:

By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

- **Do not use Version or Cipher Suite rule conditions**:



Important *Never* use either **Cipher Suite** or **Version** rule conditions in a rule with a **Decrypt - Resign** or **Decrypt - Known Key** rule action. The use of these conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

- Decrypt - Resign rule action and a Certificate Signing Request: To use a **Decrypt - Resign** rule action, you should create a Certificate Signing Request (CSR) and have it signed by a trusted certificate authority. (You can use the FMC to create a CSR: **Objects** > **Object Management** > **PKI** > **Internal Certs**.)

  To be used in a **Decrypt - Resign** rule, your certificate authority (CA) must have at least one of these extensions

  - **CA: TRUE**

    For more information, see the discussion of Basic Constraints in RFC3280, section 4.2.1.10.

  - **KeyUsage=CertSign**

    For more information see RFC 5280, section 4.2.1.3.

  To verify your CSR or CA has at least one of the preceding extensions, you can use the **openssl** command as discussed in a reference such as the openssl documentation.

This is necessary because for **Decrypt - Resign** inspection to work, the certificate that used in the decryption policy generates certificates on-the-fly and signs them so as to act as man-in-the middle and proxy all TLS/SSL connections.

- **Certificate pinning**: If the customer's browser uses certificate pinning to verify a server certificate, you cannot decrypt this traffic by re-signing the server certificate. To allow this traffic, configure a decryption rule with the **Do Not Decrypt** action to match the server certificate common name or distinguished name.
- **Non-matching cipher suite**: This error is displayed if you attempt to save a decryption rule with a cipher suite that does not match the certificate.

  ```
  Traffic cannot match this rule; none of your selected cipher suites contain a
  signature algorithm that the resigning CA's signature algorithm
  ```

- **Untrusted Certificate Authority**: If the client does not trust the Certificate Authority (CA) used to re-sign the server certificate, it warns the user that the certificate should not be trusted. To prevent this, import the CA certificate into the client trusted CA store. Alternatively, if your organization has a private PKI, you can issue an intermediate CA certificate signed by the root CA which is automatically trusted by all clients in the organization, then upload that CA certificate to the device.
- **HTTP proxy limitation**: The system cannot decrypt traffic if an HTTP proxy is positioned between a client and your managed device, and the client and server establish a tunneled TLS/SSL connection using the CONNECT HTTP method. The **Handshake Errors** undecryptable action determines how the system handles this traffic.
- **Upload signed CA**: If you create an internal CA object and choose to generate a certificate signing request (CSR), you cannot use this CA for a **Decrypt - Resign** action until you upload the signed certificate to the object.
- **Mismatched signature algorithm**: If you configure a rule with the **Decrypt - Resign** action, and mismatch signature algorithm type for one or more external certificate objects or cipher suites, the policy editor displays an **Information** ( ) next to the rule. If you mismatch signature algorithm type for all external certificate objects, or all cipher suites, the policy displays a warning icon **Warning** ( ) next to the rule, and you cannot deploy the access control policy associated with the decryption policy.

TLS/SSL block guidelines

If your decryption rule is associated with an access control policy that has a rule with an action of either **Interactive Block** or **Interactive Block with reset**, the system displays a customizable response page.

If you enabled logging in your rule, two connection events are displayed (in **Analysis** > **Unified Events**): One event for the interactive block and another event to indicate whether or not the user chose to continue to the site or not.

Related Topics

HTTP response pages, on page 2103

TLS/SSL certificate pinning guidelines

Some applications use a technique referred to as *TLS/SSL pinning* or *certificate pinning*, which embeds the fingerprint of the original server certificate in the application itself. As a result, if you configured a decryption rule with a **Decrypt - Resign** action, when the application receives a resigned certificate from a managed device, validation fails and the connection is aborted.

Because TLS/SSL pinning is used to avoid man-in-the-middle attacks, there is no way to prevent or work around it. We recommend adding a **Do Not Decrypt** rule before the **Decrypt - Resign** rule so pinning traffic is excluded from being decrypted.

- Create a **Do Not Decrypt** for those applications rule ordered before **Decrypt - Resign** rules.
- Instruct users to access the applications using a web browser.

For more information about rule ordering, see Best practice for Decryption policy rule order, on page 2196.

To determine whether applications are using TLS/SSL pinning, see Troubleshoot TLS/SSL pinning, on page 2231.

TLS/SSL heartbeat guidelines

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by RFC6520. TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

You can configure a **Max Heartbeat Length** in a Network Analysis Policy (NAP) to determine how to handle TLS heartbeats. For more information, refer to *The SSL Preprocessor* section in the "Custom Snort 2 Intrusion Policies for Access Control" document.

For more information, refer to TLS heartbeat, on page 2229.

TLS/SSL anonymous cipher suite limitation

Anonymous cipher suites impact decryption rules.

By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

You can add an anonymous cipher suite to the **Cipher Suite** condition in a decryption rule, but the system automatically strips anonymous cipher suites during ClientHello processing. For the system to use the rule, you must also configure your decryption rules in an order that prevents ClientHello processing. For more information, refer to Best practice for Decryption policy rule order, on page 2196

TLS/SSL normalizer guidelines

If you enable the **Normalize Excess Payload** option in the inline normalization preprocessor, when the preprocessor normalizes decrypted traffic, it might drop a packet and replace it with a trimmed packet. This does not end the TLS/SSL session. If the traffic is allowed, the trimmed packet is encrypted as part of the TLS/SSL session.

Other Rule-based decryption rule guidelines

Consider users and groups configuration when creating decryption rules. If you add a group or user to a rule, then change your realm settings to exclude that group or user, the rule has no effect. (The same applies to

disabling the realm.) For more information about realms, see Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

Consider all relevant categories when configuring decryption rules. If your decryption policy has a **Decrypt - Resign** action but web sites are not being decrypted, check **Category** page on rules associated with that policy.

In some cases, a web site redirects to another site for authentication or other purposes and the redirected site might have a different URL categorization than the site you're trying to decrypt. For example, `gmail.com` (**Web based email** category) redirects to `accounts.gmail.com` (**Internet Portals** category) for authentication. Be sure to include all relevant categories in the SSL rule.

Note In order to fully process traffic based on URL category, you must also configure URL filtering. See the URL Filtering Rules, on page 2085 chapter.

Consider configuring cloud query settings for URLs not in the local database. If you create a **Decrypt - Resign** rule and users browse to a web site whose category and reputation are not in the local database, data might not be decrypted. Some web sites are not categorized in the local database and, if not, data from those web sites is not decrypted by default.

You can control this behavior with the **Query Cisco cloud for unknown URLs** setting in **Integration** > **Other Integrations** > **Cloud Services**

Best practice for Decryption policy rule order

General rule ordering

Order your rules with specific conditions (such as IP addresses and networks) *before* rules with general conditions (such as applications).

Allow traffic from certificate pinned sites

Some applications use a technique referred to as *TLS/SSL pinning* or *certificate pinning*, which embeds the fingerprint of the original server certificate in the application itself. As a result, if you configured a decryption rule with a **Decrypt - Resign** action, when the application receives a resigned certificate from a managed device, validation fails and the connection is aborted.

- To confirm that TLS/SSL pinning is occurring, attempt to log in to a mobile application like Facebook. If a network connection error is displayed, log in using a web browser. (For example, you *cannot* log in to a Facebook mobile application but *can* log in to Facebook using Safari or Chrome.) You can use Firepower Management Center connection events as further proof of TLS/SSL pinning
- Configure an SSL rule with the **Do Not Decrypt** action to match the server certificate common name or distinguished name. In the SSL policy, order this rule before all **Decrypt - Resign** rules that also match the traffic.

You can retrieve the pinned certificate from the client's browser after a successful connection to the website. You can also view the certificate from the logged connection event, regardless of whether the connection succeeded or failed.

Note TLS/SSL pinning is not limited to mobile applications.

Prioritize ClientHello modifications

Place rules that match on conditions that are available in the ClientHello message before rules that match on ServerHello or server Certificate conditions.

When a managed device processes an SSL handshake, it can modify the ClientHello message to increase the likelihood of decryption. For example, it may remove compression methods because the system cannot decrypt compressed sessions.

The system modifies ClientHello messages only if it can conclusively match them to an SSL rule with either a **Decrypt - Resign**or **Decrypt - Known Key**action. The first time the system detects an encrypted session to a new server, server Certificate data is not available for ClientHello processing, which can result in an undecrypted first session. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with server Certificate conditions and process the message to maximize decryption potential.

If you place rules that match on ServerHello or server Certificate conditions (certificate, distinguished names, certificate status, cipher suites, version) before rules that match on ClientHello conditions (zones, networks, VLAN tags, ports, users, applications, URL categories), you can preempt ClientHello modification and increase the number of undecrypted sessions.

SSL policy bypass situations

The SSL policy is bypassed for any connections that match access control rules with actions of **Trust**, **Block**, or **Block with reset** if those rules:

- Use security zone, network, geolocation, and port only as the traffic matching criteria.
- Precede other rules that require inspection, such as rules that match connections based on application or URL, or allow rules that apply intrusion or file inspection.

Rule-based decryption rule traffic handling

Rule-based decryption rule traffic handling is a security mechanism that matches traffic to rule-based decryption rules in a specified order to determine appropriate actions.

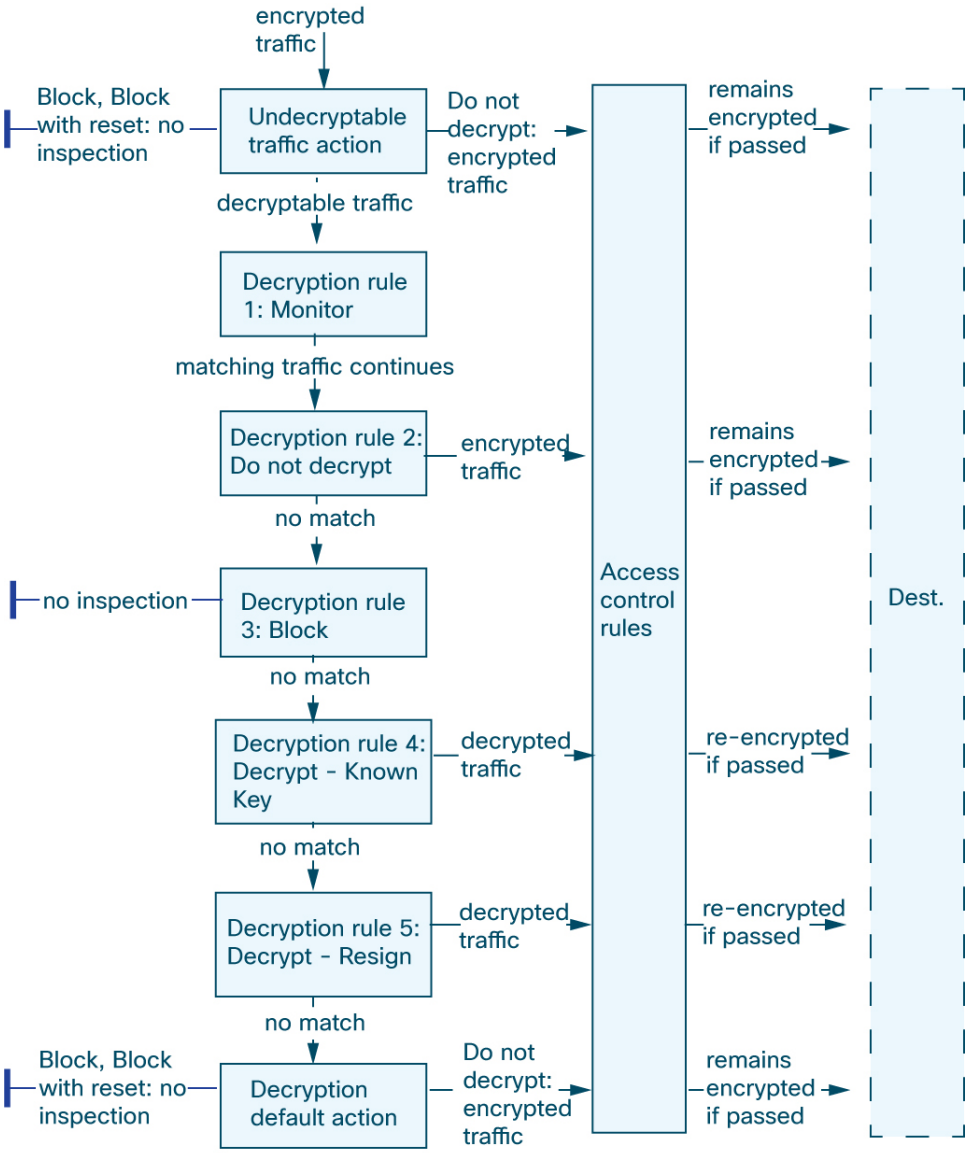
Traffic matching and processing behavior

The system matches traffic to rule-based decryption rules in the order you specify. In most cases, the system handles encrypted traffic according to the *first* decryption rule where *all* the rule's conditions match the traffic. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, VLAN, port, application, requested URL, user, certificate, certificate distinguished name, certificate status, cipher suite, or encryption protocol version.

Each rule also has an *action*, which determines whether you monitor, block, or inspect matching encrypted or decrypted traffic with access control. Note that the system does *not* further inspect encrypted traffic it blocks. It does inspect encrypted and undecryptable traffic with access control. However, some access control rule conditions require unencrypted traffic, so encrypted traffic may match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads.

Traffic handling process

This scenario summarizes the ways that Rule-based decryption rule handle traffic in an inline deployment.



In this scenario, traffic is evaluated as follows:

- Undecryptable Traffic Action evaluates encrypted traffic first. For traffic the system cannot decrypt, the system either blocks it without further inspection or passes it for access control inspection. Encrypted traffic that does not match continues to the next rule.
- Rule 1: Monitor evaluates encrypted traffic next. Monitor rules track and log encrypted traffic but do not affect traffic flow. The system continues to match traffic against additional rules to determine whether to permit or deny it.

- Rule 2: Do Not Decrypt evaluates encrypted traffic third. Matching traffic is not decrypted; the system inspects this traffic with access control, but not file or intrusion inspection. Traffic that does not match continues to the next rule.
- Rule 3: Block evaluates encrypted traffic fourth. Matching traffic is blocked without further inspection. Traffic that does not match continues to the next rule.
- Rule 4: Decrypt - Known Key evaluates encrypted traffic fifth. Matching traffic incoming to your network is decrypted using a private key you upload. The decrypted traffic is then evaluated against access control rules. Access control rules handle decrypted and unencrypted traffic identically. The system can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the decryption rule continues to the next rule.
- Rule 5: Decrypt - Resign is the final rule. If traffic matches, the system re-signs the server certificate with an uploaded CA certificate and decrypts traffic. The decrypted traffic is then evaluated against access control rules, which treat decrypted and unencrypted traffic the same. This additional inspection allows the system to block traffic. All remaining traffic is re-encrypted before being sent to its destination. Traffic that does not match the decryption rule rule continues to the next rule.
- Default Action handles all traffic that does not match any of the decryption rule. The default action either blocks encrypted traffic without further inspection or does not decrypt it, passing it for access control inspection.

Encrypted traffic inspection configuration

You must create reusable public key infrastructure (PKI) objects to control encrypted traffic based on encrypted session characteristics and decrypt encrypted traffic. You can add this information on the fly when uploading trusted certificate authority (CA) certificates to the a decryption policy and creating decryption rule, creating the associated object in the process. However, configuring these objects ahead of time reduces the chance of improper object creation.

Decrypting encrypted traffic with certificates and paired keys

The system can decrypt incoming encrypted traffic if you configure an internal certificate object by uploading the server certificate and private key used to encrypt the session. If you reference that object in a decryption policy rule with an action of **Decrypt - Known Key** and traffic matches that rule, the system uses the uploaded private key to decrypt the session.

The system can also decrypt outgoing traffic if you configure an internal CA object by uploading a CA certificate and private key. If you reference that object in a decryption rule with an action of **Decrypt - Resign** and traffic matches that rule, the system re-signs the server certificate passed to the client browser, then acts as a man-in-the-middle to decrypt the session. You can optionally replace the self-signed certificate key only and not the entire certificate, in which case users see a self-signed certificate key notice in the browser.

Controlling traffic based on encrypted session characteristics

The system can control encrypted traffic based on the cipher suite or server certificate used to negotiate the session. You can configure one of several different reusable objects and reference the object in a decryption rule condition to match traffic. This table describes the different types of reusable objects you can configure:

If you configure...	You can control encrypted traffic based on whether...
A cipher suite list containing one or more cipher suites	The cipher suite used to negotiate the encrypted session matches a cipher suite in the cipher suite list
A trusted CA object by uploading a CA certificate your organization trusts	The trusted CA trusts the server certificate used to encrypt the session, whether: • The CA issued the certificate directly • The CA issued a certificate to an intermediate CA that issued the server certificate
An external certificate object by uploading a server certificate	The server certificate used to encrypt the session matches the uploaded server certificate
A distinguished name object containing a certificate subject or issuer distinguished name	The subject or issuer common name, country, organization, or organizational unit on the certificate used to encrypt the session matches the configured distinguished name

Related Topics

Cipher suite lists, on page 2958
Distinguished names, on page 2962
PKI objects, on page 2986

Rule-based decryption rule order evaluation

Rule-based decryption rule order evaluation determines the sequence in which the system evaluates decryption rules in a policy.

Rule evaluation process

This information applies only to rule-based decryption policies and rules.

When you create a rule-based decryption rule in a rule-based decryption policy , you specify its position using the **Insert** list in the rule editor. Decryption rules in an a decryption policy are numbered, starting at 1. The system matches traffic to decryption rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* decryption rule where *all* the rule's conditions match the traffic. Except in the case of Monitor rules (which log traffic but do not affect traffic flow), the system does *not* continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, VLAN, port, application, requested URL category and reputation, user, certificate, certificate distinguished name, certificate status, cipher suite, or encryption protocol version.

Each rule also has an *action*, which determines whether you monitor, block, or inspect matching encrypted or decrypted traffic with access control. Note that the system does *not* further inspect encrypted traffic it blocks. It does subject encrypted and undecryptable traffic to access control. However, access control rule conditions require unencrypted traffic, so encrypted traffic matches fewer rules.

Rules that use *specific* conditions (such as network and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network)

should be ordered first in your rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your rules.



Note Proper decryption rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

In addition to ordering rules by number, you can group rules by category. By default the system provides three categories: Administrator, Standard, and Root. You can add custom categories, but you cannot delete the system-provided categories or change their order.

Related Topics

Best practices for access control rules, on page 2058

Rule-based decryption rule conditions

A rule-based decryption rule condition identifies the type of encrypted traffic the rule handles. Conditions can be simple or complex, and you can specify more than one condition type per rule. Only if traffic meets all the conditions in a rule does the rule apply to the traffic.

Condition evaluation and rule actions

If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion. For example, a rule with a certificate condition but no version condition evaluates traffic based on the server certificate used to negotiate the session, regardless of the session SSL or TLS version.

Every rule-based decryption policy has an associated action that determines handling and logging for matching encrypted traffic:

- Handling: Most importantly, the rule action governs whether the system will monitor, trust, block, or decrypt encrypted traffic that matches the rule's conditions.
- Logging: The rule action determines when and how you can log details about matching encrypted traffic.

Your TLS/SSL inspection configuration handles, inspects, and logs decrypted traffic:

- The decryption policy's undecryptable actions handle traffic that the system cannot decrypt.
- The policy's default action handles traffic that does not meet the condition of any non-Monitor decryption rule.

You can log a connection event when the system blocks or trusts an encrypted session. You can also force the system to log connections that it decrypts for further evaluation by access control rules, regardless of how the system later handles or inspects the traffic. Connection logs for encrypted sessions contain details, such as the certificate used to encrypt that session. You can log only end-of-connection events.

- For blocked connections (Block, Block with reset), the system immediately ends the sessions and generates an event.
- For Do Not Decrypt connections, the system generates an event when the session ends.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Caution Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include a decryption policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

Related Topics

Security zone rule conditions, on page 2202
Network rule conditions, on page 1701
VLAN tag rule conditions, on page 2070
User rule conditions, on page 1701
Application rule conditions, on page 1701
Best practice: use application-based rules over port-based rules, on page 1703
Category rule conditions, on page 2206
Server certificate-based rule conditions, on page 2207

Security zone rule conditions

A security zone rule condition is a network traffic control mechanism that

- segments your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices,
- controls or decrypts traffic by its source and destination security zones, and
- requires all interfaces in a zone to be of the same type (all inline, passive, switched, or routed).

Zone condition traffic matching

If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

All zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Network rule conditions

A network condition is a traffic matching criterion that

- controls or decrypts traffic based on source and destination IP addresses using inner packet headers,
- enables administrators to specify exact IP addresses or blocks for granular policy enforcement, and
- can be built using predefined network objects or by manually entering IP addresses or address blocks.

Network rule condition restrictions

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

VLAN tag rule conditions

VLAN tags rule conditions are filtering mechanisms that

- control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic
- use the innermost VLAN tag to filter VLAN traffic, with the exception of prefilter policy which uses the outermost VLAN tag, and
- apply only to inline sets and do not match traffic on firewall interfaces.

VLAN tag rule conditions configuration details

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note this Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).

- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

User rule conditions

A user rule condition is a traffic matching mechanism that

- matches traffic based on whether the user initiating the connection or the group to which they belong
- can be configured for users in Microsoft Active Directory realms only, and
- enables policy enforcement based on user identity and group membership.

Special identities users

In addition to configuring users and groups for configured realms, you can set policies for the special identities users:

- Failed Authentication: User who failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.
- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified such as users who are not downloaded by a configured realm.

For access control rules only, you must first associate an identity policy with the access control policy as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

User rule condition example

Configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

Application rule conditions

Application rule conditions are application awareness and control mechanisms that

- identify and classify applications through IP traffic analysis
- enable control with organized filters
- require one enabled detector for each condition.

Application filters and control benefits

System-provided *application filters* help perform control by organizing applications according to basic characteristics, such as type, risk, business relevance, category, and tags. You can also create reuseable, user-defined filters using combinations of system-provided filters, or custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled, the system automatically enables all available system-provided detectors, or the most recently modified user-defined detector. For more information about application detectors, see Application detectors.

You can utilize both application filters and individually specified applications to ensure complete coverage. Understand the application's requirements to effectively order access control rules.

Application filters help you to quickly configure application control effectively. For example, you can easily use system-provided filters to create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

The system characterizes detected applications using several criteria which form effective filters.

Table 125: Application characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email applications are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Best practice: use application-based rules over port-based rules

Use application filtering criteria instead of port filtering to target traffic. Applications can be configured to use unique ports to bypass access control blocks.

Port conditions allow you to control traffic by its source and destination ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

The traditional method targets applications by using port specifications, but applications can manipulate unique ports to bypass these controls

Application filtering is recommended for applications to avoid blocking legitimate applications, like Firewall Threat Defense, that dynamically open separate channels for control and data flow. Port-based access control rules could block such applications, impacting desirable operations.

Using source and destination port constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP).

- For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).
- If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Category rule conditions

Category rule conditions are policy elements that

- enable control of web access based on Cisco Talos intelligence
- support policy creation through URL filtering without relying on company vertical, intent, or declared security, and
- can be included in decryption policies to manage encrypted traffic.

Additional information

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security.

For more information, see URL filtering, on page 2085.

If you are using category rule conditions in decryption policies in a rule with the **Do Not Decrypt** rule action, see Decryption rule Do Not Decrypt action, on page 2221.

Server certificate-based rule conditions

Server certificate-based rule conditions are decryption rules configurations that

- handle and decrypt encrypted traffic based on server certificate characteristics
- inspect traffic using certificate attributes, status, cipher suites, and session versions, and
- enhance secure traffic management and visibility through certificate validation criteria.

Server certificate attribute conditions

You can configure decryption rules based on these server certificate attributes:

- **Distinguished name conditions**: Allow you to handle and inspect encrypted traffic based on the CA that issued a server certificate, or the certificate holder. Based on the issuer distinguished name, you can handle traffic based on the CA that issued a site's server certificate.
- **Certificate conditions**: In decryption rules allow you to handle and inspect encrypted traffic based on the server certificate used to encrypt that traffic. You can configure a condition with one or more certificates; traffic matches the rule if the certificate matches any of the condition's certificates.
- **Certificate status conditions**: In decryption rules allow you to handle and inspect encrypted traffic based on the status of the server certificate used to encrypt the traffic, including whether a certificate is valid, revoked, expired, not yet valid, self-signed, signed by a trusted CA, whether the Certificate Revocation List (CRL) is valid; whether the Server Name Indication (SNI) in the certificate matches the server in the request.
- **Cipher suite conditions**: In decryption rules allow you to handle and inspect encrypted traffic based on the cipher suite used to negotiate the encrypted session.
- **Session conditions**: In decryption rules allow you to inspect encrypted traffic based on the SSL or TLS version used to encrypt the traffic.

To detect multiple cipher suites in a rule, the certificate issuer, or the certificate holder, you can create reusable cipher suite list and distinguished name objects and add them to your rule. To detect the server certificate and certain certificate statuses, you must create external certificate and external CA objects for the rule.

Related Topics

Certificate rule conditions, on page 2207
Certificate Status Decryption rule conditions, on page 2214
Trust external certificate authorities, on page 2213
Match traffic on certificate status
Cipher suite rule conditions, on page 2217
Encryption protocol version rule conditions, on page 2220

Certificate rule conditions

When you build a certificate-based decryption rule condition, you can upload a server certificate and save it as an external certificate *object* The certificate object is reusable and associates a name with a server certificate. Alternately, you can configure certificate conditions with existing external certificate objects and object groups.

You can search the **Available Certificates** field in the rule condition based for external certificate objects and object groups based on these certificate distinguished name characteristics:

- Subject or issuer common name (CN), or if the URL is contained in the certificate's Subject Alternative Name (SAN)

  The URL the user enters in the browser matches the Common Name (CN)

- Subject or issuer organization (O)
- Subject or issuer organizational unit (OU)

You can choose to match against multiple certificates in a single certificate rule condition. If the certificate used to encrypt the traffic matches any of the uploaded certificates, the encrypted traffic matches the rule.

You can add a maximum of 50 external certificate objects and external certificate object groups to the **Selected Certificates** in a single certificate condition.

Note these items:

- You cannot configure a certificate condition if you also select the **Decrypt - Known Key** action. That action requires you to select a server certificate to decrypt traffic. Therefore, the certificate already matches the traffic.
- If you configure a certificate condition with an external certificate object, any cipher suites you add to a cipher suite condition or internal CA objects you associate with the **Decrypt - Resign** action, must match the external certificate's signature algorithm type. For example, if your rule's certificate condition references an EC-based server certificate, any cipher suites you add or CA certificates you associate with the **Decrypt - Resign** action, must also be EC-based. If you mismatch signature algorithm types in this case, the policy editor displays a warning next to the rule.
- When the system first detects an encrypted session to a new server, it does not have certificate data for ClientHello processing. As a result, the first session may not be decrypted. After the initial session, the managed device caches data from the server Certificate message. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with certificate conditions and process the message to maximize decryption potential.

Conditions for distinguished name rules in decryption policies

Learn to use distinguished name conditions in a decryption rule. If you're not sure, you can find a certificate's Subject Alternative Name (SAN) and Common Name using a web browser. You can then you can add those values to a decryption rule as distinguished name conditions.

For detailed information about SANs, see RFC 5280, section 4.2.1.6.

Overview

Distinguished name (DN) conditions allow you to specify certificate field values in decryption rules. This helps control which encrypted traffic is decrypted or excluded from decryption based on the certificate’s Common Name (CN) and Subject Alternative Name (SAN)

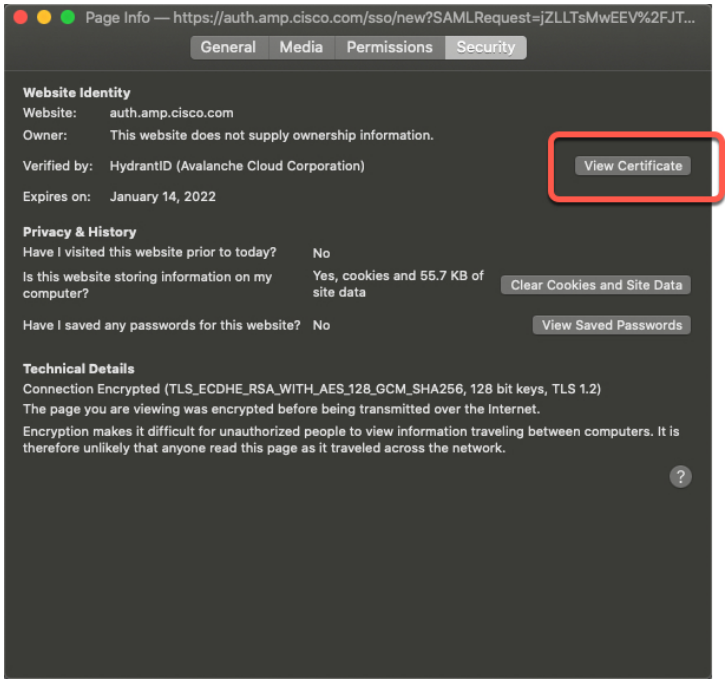
Identifying Certificate Values

Use these steps to find a certificate's Common Name. These steps also work to find the Common Name and SANs for a self-signed certificate.

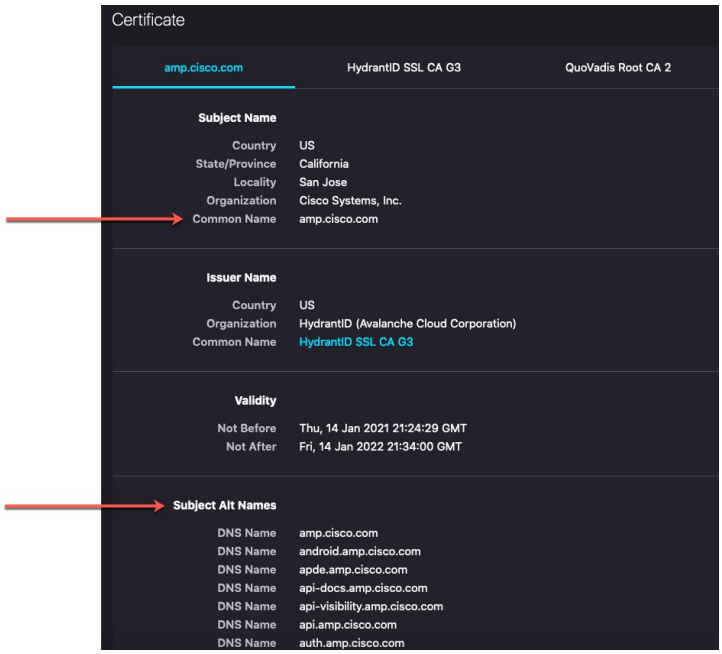
These instructions use Firefox as an example, but other browsers work similarly. The following procedure uses `amp.cisco.com` as an example.

1. Browse to `amp.cisco.com` in Firefox.
2. In the browser's location bar, to the left of the URL, click .
3. Click **Connection secure** > **More Information**.

   (For a non-secure or self-signed certificate, click **Connection not secure** > **More Information**.)
4. On the Page Info dialog box, click **View Certificate**.



5. The next page shows certificate details.



Note the following:

- `CN=auth.amp.cisco.com`, if used as a DN rule condition, would match ONLY that host name (that is, SNI). The SNI `amp.cisco.com` would NOT match.
- Wildcards can be used in DNs to match multiple domains.

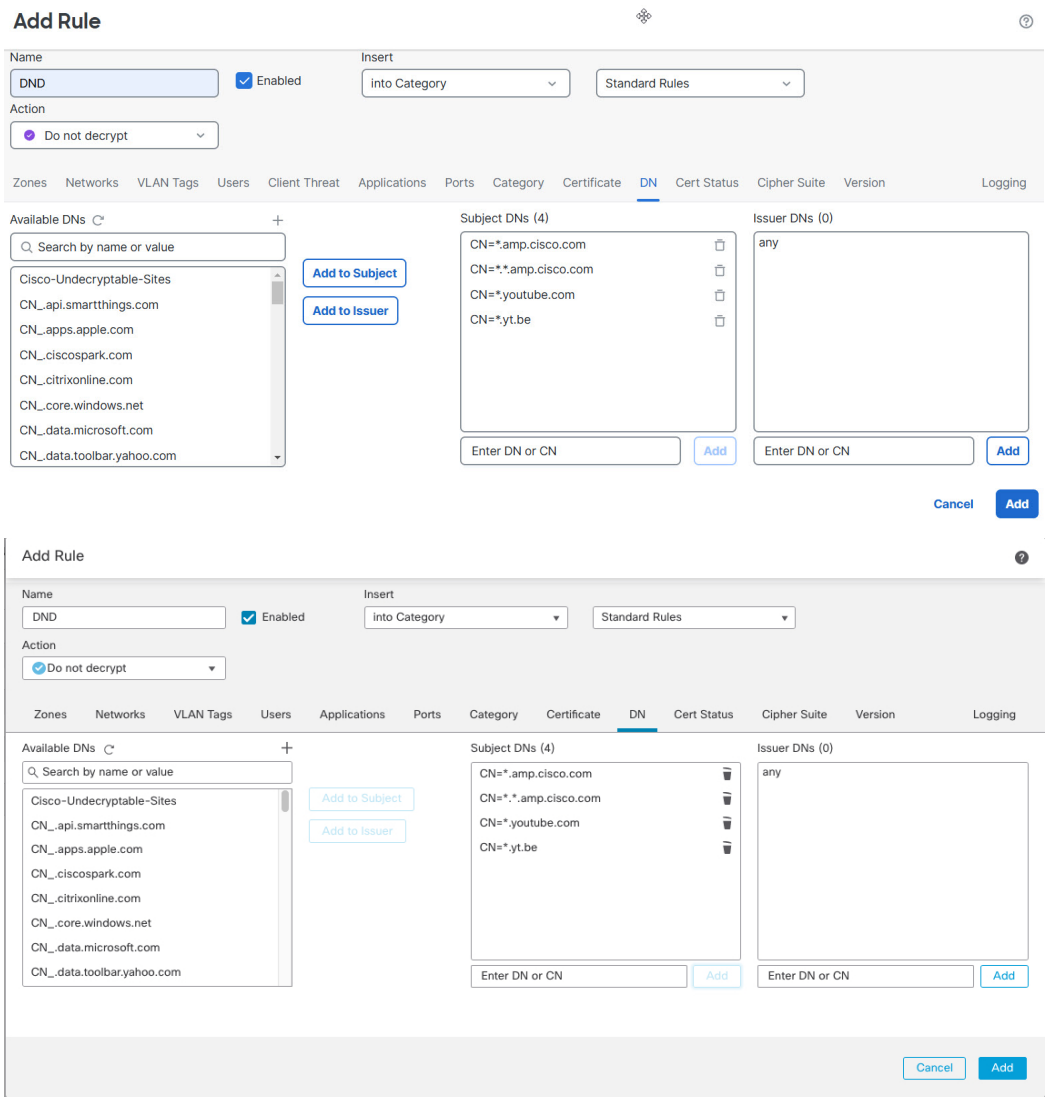
  For example, to match `auth.amp.cisco.com`, use `CN=*.amp.cisco.com`. To match `auth.us.amp.cisco.com`, use `CN=*.*.amp.cisco.com`.

  A DN like `CN=*.example.com` matches `www.example.com` but NOT `example.com`. To match both SNIs, use two DNs in the rule condition.
- Avoid overly broad wildcards like `CN=*.google.com` matches a very large number of SANs. Instead of `CN=*.google.com`, use a DN object like `CN=*.youtube.com` as the DN object so it matches names like `www.youtube.com`.

  You can also use variations of the SNI that match SANs like `CN=*.youtube.com`, `CN=youtu.be`, `CN=*.yt.be`, and so on.
- A self-signed certificate should work the same way. You can confirm it's a self-signed certificate by the fact the issuer DN is the same as the subject DN.

Example: DN rule matching

Suppose you want to prevent decryption for traffic going to `amp.cisco.com` or YouTube. Configure DN rule conditions such as:



These DN conditions would match URLs like:

- www.amp.cisco.com
- auth.amp.cisco.com
- auth.us.amp.cisco.com
- www.youtube.com
- kids.youtube.com
- www.yt.be

These DN conditions would not match any of the URLs listed, so the traffic would not match the Do Not Decrypt rule. However, it might match another decryption rules in the same decryption policy .

- amp.cisco.com
- youtube.com

- yt.be

To match any of the preceding host names, add more CNs to the rule (for example, adding `CN=yt.be` would match that URL.)

System behavior: SNI and SANs usage

The host name in the client request is called the Server Name Indication (SNI). During the TLS handshake, the client specifies which hostname they want to connect to (for example, `auth.amp.cisco.com`) using the SNI extension. The server then selects the corresponding private key and certificate chain that are required to establish the connection while hosting all certificates on a single IP address.

If the SNI matches the CN or a SAN in the certificate, the SNI is used for DN rule matching.

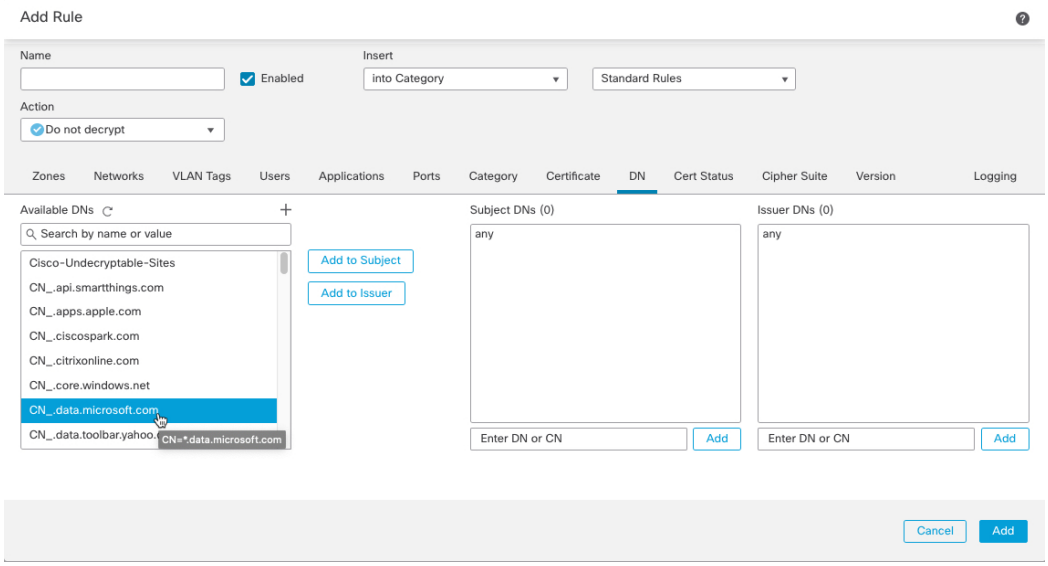
If there is no SNI, or the SNI does not match the certificate, the CN is used for DN rule matching.

Add a DN rule condition

After you identify the CN you want to match,, edit the decryption rule by using an existing DN, creating a new DN object, or adding the DN directly.

- Use an existing DN.

  Click the name of a DN and then click either **Add to Subject** or **Add to Issuer**. (**Add to Subject** is much more common.) To view the value of a DN object, hover the mouse pointer over it.)

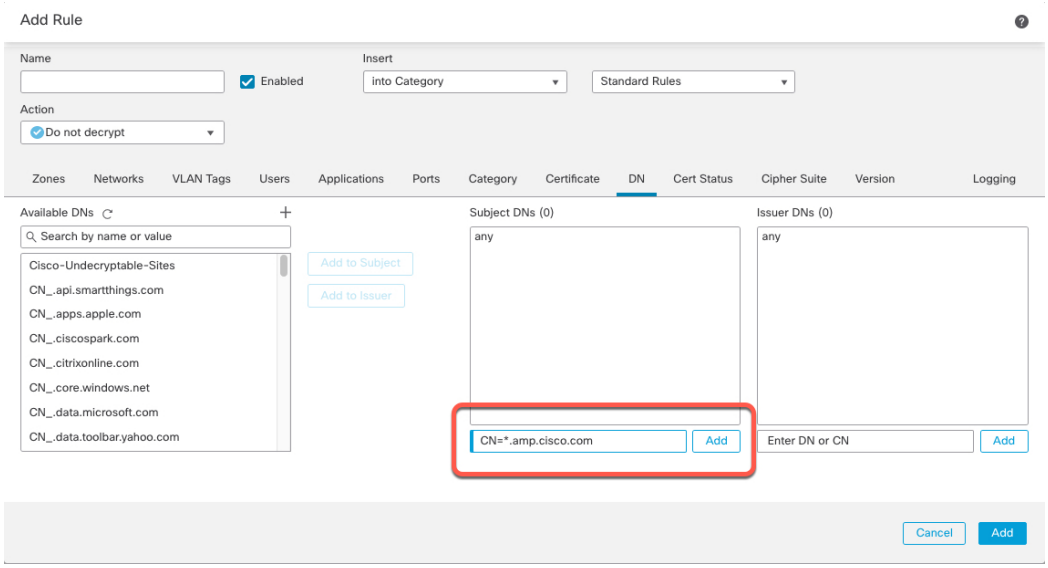


- Create a new DN object.

  Click **Add** (+) to the right of Available DNs. The DN object must consist of a name and a value.

- Add the DN directly.

  Enter the DN in the field at the bottom of the **Subject DNs** field or the **Issuer DNs** field. (**Subject DNs** is more common.) After you enter the DN, click **Add**.



Related Topics

Distinguished names, on page 2962

Trust external certificate authorities

Trust CAs by adding root and intermediate CA certificates to your decryption policy, then use these trusted CAs to verify server certificates used to encrypt traffic.

If a trusted CA certificate contains an uploaded certificate revocation list (CRL), you can also verify whether a trusted CA revoked the encryption certificate.

Tip Upload all certificates in a root CA's chain of trust to the list of trusted CA certificates, including the root CA certificate and all intermediate CA certificates. Otherwise, it is more difficult to detect trusted certificates issued by intermediate CAs. Also, if you configure certificate status conditions to trust traffic based on the root issuer CA, all traffic within a trusted CA's chain of trust can be allowed without decryption, rather than unnecessarily decrypting it.

For more information, see Trusted CA objects, on page 2993.

Note When you create a decryption policy, the policy's **Trusted CA Certificate** tab page is populated with several trusted CA certificates, including the **Cisco-Trusted-Authorities** group, which is added to the **Select Trusted CAs** list.

Procedure

Step 1 Click **Policies** > **Access Control heading** > **Decryption**.

Step 2 Click **Edit** ( ) next to the decryption policy to edit.

Step 3 Click **Add Rule** to add a new decryption rule or click **Edit** ( ) to edit an existing rule.

Step 4 Click the **Certificates** tab.

Step 5 Find the trusted CAs you want to add from the **Available Certificates**, as follows:

- To add a trusted CA object on the fly, which you can then add to the condition, click **Add** ( ) above the **Available Certificates** list.
- To search for trusted CA objects and groups to add, click the **Search by name or value** prompt above the **Available Certificates** list, then enter either the name of the object, or a value in the object. The list updates as you type to display matching objects.

Step 6 To select an object, click it. To select all objects, right-click and then **Select All**.

Step 7 Click **Add to Rule**.

Tip
You can also drag and drop selected objects.

Step 8 Add or continue editing the rule.

The trusted CA certificates are configured in your decryption policy and can be used to verify server certificates for secure traffic inspection.

What to do next

- Add a certificate status decryption rule condition to your SSL rule.
- Deploy configuration changes.

Certificate Status Decryption rule conditions

For each certificate status decryption rule you configure, you can match traffic against the presence or absence of a given status. You can select several statuses in one rule condition. If the certificate matches any of the selected statuses, the rule matches the traffic.

You can choose to match against the presence or absence of multiple certificate statuses in a single certificate status rule condition. The certificate needs to match only one of the criteria to match the rule.

You should consider, when setting this parameter, whether you're configuring a decrypt rule or a block rule. Typically, you should click **Yes** for a block rule and **No** for a decrypt rule. Examples:

- If you're configuring a **Decrypt - Resign** rule, the default behavior is to decrypt traffic with an expired certificate. To change that behavior, click **No** for **Expired** so traffic with an expired certificate is not decrypted and resigned.
- If you're configuring a **Block** rule, the default behavior is to allow traffic with an expired certificate. To change that behavior click **Yes** for **Expired** so traffic with an expired certificate is blocked.

This table describes how the system evaluates encrypted traffic based on the status of the encrypting server's certificate.

Table 126: Certificate status rule condition criteria

Status Check	Status Set to Yes	Status Set to No
Revoked	The policy trusts the CA that issued the server certificate, and the CA certificate uploaded to the policy contains a CRL that revokes the server certificate.	The policy trusts the CA that issued the certificate, and the CA certificate uploa policy does not contain a CRL that rev certificate.
Self-signed	The detected server certificate contains the same subject and issuer distinguished name.	The detected server certificate contains subject and issuer distinguished names
Valid	All of these are true: • The policy trusts the CA that issued the certificate. • The signature is valid. • The issuer is valid. • None of the policy's trusted CAs revoked the certificate. • The current date is between the certificate Valid From and Valid To date.	At least one of these is true: • The policy does not trust the CA t certificate. • The signature is invalid. • The issuer is invalid. • A trusted CA in the policy revoke certificate. • The current date is before the cert From date. • The current date is after the certifi date.
Invalid signature	The certificate's signature cannot be properly validated against the certificate's content.	The certificate's signature is properly val the certificate's content.
Invalid issuer	The issuer CA certificate is not stored in the policy's list of trusted CA certificates.	The issuer CA certificate is stored in th of trusted CA certificates.
Expired	The current date is after the certificate Valid To date.	The current date is before or on the cer To date.
Not yet valid	The current date is before the certificate Valid From date.	The current date is after or on the certi From date.

Status Check	Status Set to Yes	Status Set to No
Invalid certificate	The certificate is not valid. At least one of these is true: • Invalid or inconsistent certificate extension; that is, a certificate extension had an invalid value (for example, an incorrect encoding) or some value inconsistent with other extensions. • The certificate cannot be used for the specified purpose. • The Basic Constraints path length parameter has been exceeded. For more information, see RFC 5280, section 4.2.1.9. • The certificate's value for Not Before or Not After is invalid. These dates can be encoded as UTCTime or GeneralizedTime For more information, see RFC 5280 section 4.1.2.5. • The format of the name constraint is not recognized; for example, an email address format of a form not mentioned in RFC 5280, section 4.2.1.10. This could be caused by an improper extension or some new feature not currently supported. An unsupported name constraint type was encountered. OpenSSL currently supports only directory name, DNS name, email, and URI types. • The root certificate authority is not trusted for the specified purpose. • The root certificate authority rejects the specified purpose.	The certificate is valid. All of these are tru • Valid certificate extension. • The certificate can be used for the spe purpose. • Valid Basic Constraints path length. • Valid values for Not Before and Not A • Valid name constraint. • The root certificate is trusted for the s purpose. • The root certificate accepts the specifie

Status Check	Status Set to Yes	Status Set to No
Invalid CRL	The Certificate Revocation List (CRL) digital signature is not valid. At least one of these is true: • The value of the CRL's Next Update or Last Update field is invalid. • The CRL is not yet valid. • The CRL has expired. • An error occurred when attempting to verify the CRL path. This error occurs only if extended CRL checking is enabled. • CRL could not be found. • The only CRLs that could be found did not match the scope of the certificate.	The CRL is valid. All of these are true: • Next Update and Last Update fiel • The CRL's date is valid. • The path is valid. • The CRL was found. • The CRL matches the certificate's
Server mismatch	The server name does not match the server's Server Name Indication (SNI) name, which could indicate an attempt to spoof the server name.	The server name matches the SNI name to which the client is requesting access

Note that even though a certificate might match more than one status, the rule causes an action to be taken on the traffic only once.

Checking whether a CA issued or revoked a certificate requires uploading root and intermediate CA certificates and associated CRLs as objects. You then add these trusted CA objects to an decryption policy's list of trusted CA certificates.

Cipher suite rule conditions

The system provides predefined cipher suites you can add to a cipher suite rule condition for Block or Block WITH Reset rule actions. You can also add cipher suite list objects containing multiple cipher suites.

Important Cipher suite rule conditions should be used only to *block* traffic, they should not be used to *decrypt* traffic.

Note You cannot add new cipher suites. You can neither modify nor delete predefined cipher suites.

You can add a maximum of 50 cipher suites and cipher suite lists to the **Selected Cipher Suites** in a single cipher suite condition. The system supports adding these cipher suites to a cipher suite condition:

- SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_FIPS_WITH_DES_CBC_SHA
- TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA

- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA
- TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA
- TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256
- TLS_DHE_RSA_WITH_DES_CBC_SHA
- TLS_DH_Anon_WITH_AES_128_GCM_SHA256
- TLS_DH_Anon_WITH_AES_256_GCM_SHA384
- TLS_DH_Anon_WITH_CAMELLIA_128_CBC_SHA
- TLS_DH_anon_WITH_CAMELLIA_128_CBC_SHA256
- TLS_DH_Anon_WITH_CAMELLIA_256_CBC_SHA
- TLS_DH_anon_WITH_CAMELLIA_256_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_NULL_SHA
- TLS_ECDHE_ECDSA_WITH_RC4_128_SHA
- TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384

- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_NULL_SHA
- TLS_ECDHE_RSA_WITH_RC4_128_SHA
- TLS_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA256
- TLS_RSA_WITH_DES_CBC_SHA
- TLS_RSA_WITH_NULL_MD5
- TLS_RSA_WITH_NULL_SHA
- TLS_RSA_WITH_RC4_128_MD5
- TLS_RSA_WITH_RC4_128_SHA

Note these considerations:

- If you add cipher suites not supported for your deployment, you cannot deploy your configuration. For example, passive deployments do not support decrypting traffic with the any of the ephemeral Diffie-Hellman (DHE) or ephemeral elliptic curve Diffie-Hellman (ECDHE) cipher suites. Creating a rule with these cipher suites prevents you from deploying your access control policy.
- You can add an anonymous cipher suite to the **Cipher Suite** condition in decryption policy. To use the rule, you must also configure your settings in an order that prevents ClientHello processing. For more information, refer to Best practice for Decryption policy rule order, on page 2196.
- When you specify a cipher suite as a rule condition, the rule matches the negotiated cipher suite in the ServerHello message, not the entire list from the ClientHello message. During ClientHello processing, the managed device removes unsupported cipher suites from the ClientHello message. If all specified cipher suites are removed, the system retains the original list. If the system keeps unsupported cipher suites, it cannot decrypt the session during subsequent evaluation.

Encryption protocol version rule conditions

You can choose to match against traffic encrypted with SSL version 3.0, or TLS version 1.0, 1.1, or 1.2. By default, all protocol versions are selected when you create a rule; if you select multiple versions, encrypted traffic that matches any of the selected versions matches the rule. You must select at least one protocol version when saving the rule condition.

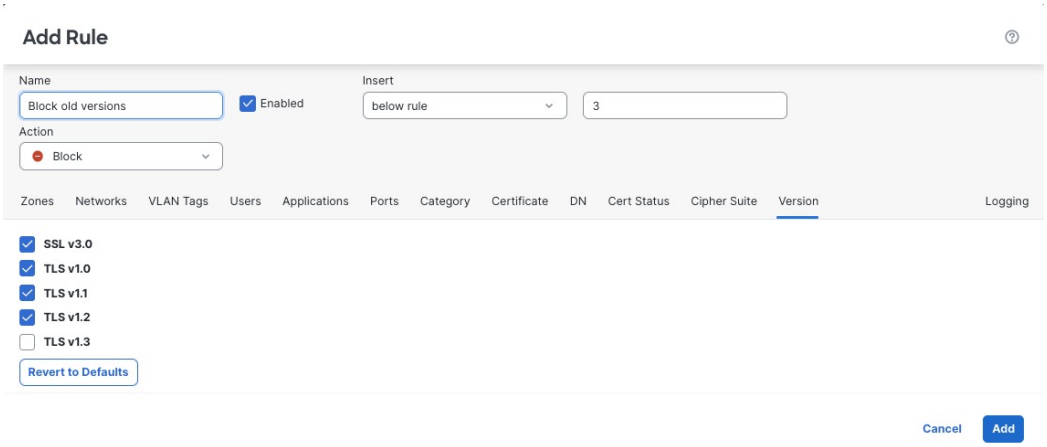
You can use SSL 3.0 in a Do Not Decrypt, Block, or Block with Reset rule action.

You *cannot* select SSL v2.0 in a version rule condition; the system does not support decrypting traffic encrypted with SSL version 2.0. You can configure an undecryptable action to allow or block this traffic without further inspection. For more information, see Set default handling for undecryptable traffic, on page 2181.

Important Protocol version rule conditions should be used *only* to *block* traffic, they should not be used to *decrypt* traffic.

For example, to block all SSL v3.0, TLS v1.0, TLS v1.1, and TLS v1.2 traffic, set the options as follows:

Figure 386: Block all SSL v3.0, TLS v1.0, TLS v1.1, and TLS 1.2 traffic



Rule-based decryption rule actions

Rule-based decryption rule actions are operational functions that define how the system processes SSL/TLS traffic based on configured decryption rules.

Available actions

The system provides various actions for rule-based decryption rule.

Decryption rule monitor action

The **Monitor** action is not designed to permit or deny traffic. Rather, its primary purpose is to force connection logging, regardless of how matching traffic is eventually handled. The ClientHello message is not modified if traffic matches a **Monitor** rule condition.

Monitor action traffic handling

Traffic that matches a **Monitor** rule is processed in this sequence:

1. Traffic matches the Monitor rule condition and connection logging is forced.
2. Traffic is then matched against additional rules, if present, to determine whether to trust, block, or decrypt it.
3. The first non-Monitor rule matched determines traffic flow and any further inspection.
4. If there are no additional matching rules, the system uses the default action.

The system automatically logs end-of connection events for monitored traffic to the Security Cloud Control database, regardless of the logging configuration of the rule or default action that later handles the connection.

Decryption rule Do Not Decrypt action

The Do Not Decrypt action passes encrypted traffic for evaluation by the access control policy's rules and default action. Because some access control rule conditions require unencrypted traffic, this traffic might match fewer rules. The system cannot perform deep inspection on encrypted traffic, such as intrusion or file inspection.

Typical use cases for do not decrypt rules

Typical reasons for a **Do Not Decrypt** rule action include:

- When decrypting TLS/SSL traffic is prohibited by law.
- Sites you know you can trust.
- Sites you can disrupt by inspecting traffic (such as Windows Update).
- To view the values of TLS/SSL fields using connection events. (You do not need to decrypt traffic to view connection event fields.) .

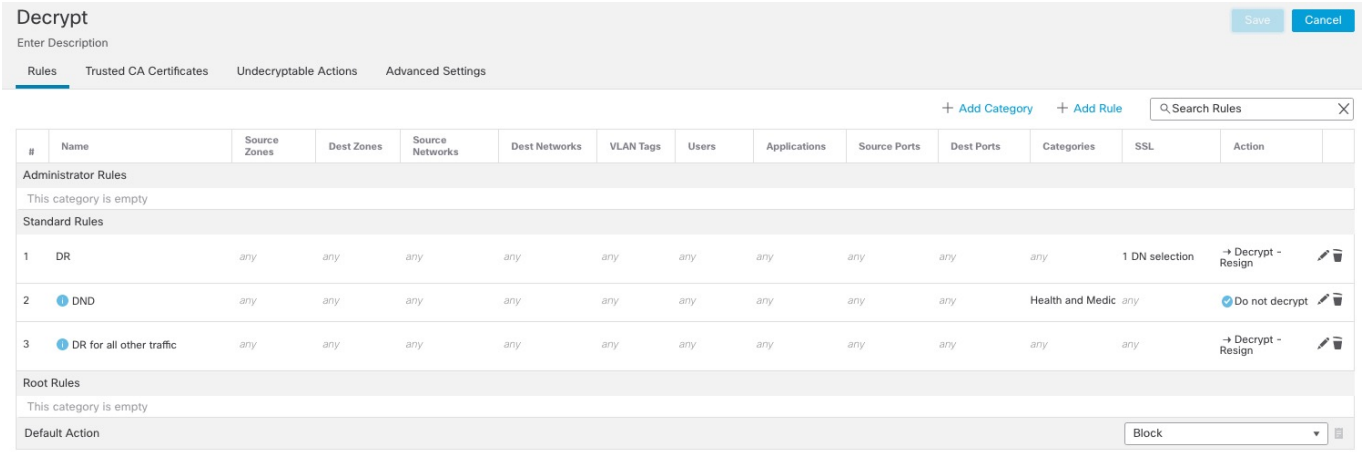
For more information, see Default handling options for undecryptable traffic, on page 2180

Limitations of categories in Do Not Decrypt rules

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security. While we strive to continuously update and improve URL filtering categories, it is not an exact science. Some websites are not categorized at all and it's possible some websites might be improperly categorized.

Avoid overusing categories in do not decrypt rules to avoid decrypting traffic without a reason; for example, the Health and Medicine category includes the WebMD website, which does not threaten patient privacy.

Following is a sample decryption policy that can prevent decryption for websites in the Health and Medicine categories but allows decryption for WebMD and everything else. General information about decryption rules can be found in Guidelines for using TLS/SSL decryption, on page 2188.



Note Don't confuse URL filtering with application detection, which relies on reading some of the packet from a website to determine more specifically what it is (for example, Facebook Message or Salesforce). For more information, see Recommendations for application control, on page 2054.

Decryption rule blocking actions

decryption rule blocking actions provide different termination methods to control connection handling.

Available blocking actions

The system provides these decryption rule actions for traffic you do not want to pass through the system:

- **Block**: Terminates the connection, resulting in an error in the client browser. The error message does not indicate the site was blocked due to policy. Instead, errors might indicate that there are no common encryption algorithms. It is not obvious from this message that you blocked the connection on purpose.
- **Block with reset**: Terminates and resets the connection, resulting in an error in the client browser. The error indicates the connection was reset but does not indicate why.

Tip You cannot use the **Block** or **Block with reset** action in a passive or inline (tap mode) deployment because the device does not directly inspect the traffic. If you create a rule with the **Block** or **Block with reset** action that contains passive or inline (tap mode) interfaces within a security zone condition, the policy editor displays a warning (⚠) next to the rule.

Decryption rule decrypt actions

Decrypt - Known Key, and **Decrypt - Resign** are rule actions that decrypt encrypted traffic. The system inspects decrypted traffic with access control. Access control rules handle decrypted and unencrypted traffic

identically — you can inspect it for discovery data as well as detect and block intrusions, prohibited files, and malware. The system reencrypts allowed traffic before passing it to its destination.

Certificate authority recommendations

We recommend you use a certificate from a trusted Certificate Authority (CA) to decrypt traffic. This prevents **Invalid Issuer** from being displayed in the SSL Certificate Status column in connection events.

For more information about adding trusted objects, see Trusted certificate authority objects, on page 2992.

Related topic: TLS 1.3 decryption best practices, on page 2184.

Related Topics

TLS 1.3 decryption best practices, on page 2184

Incoming traffic decryption actions

Both the **Decrypt - Replace Cert** and **Decrypt - Known Key** decryption actions enable you to decrypt incoming traffic to a server that you want to protect from malicious attacks.

Decryption action comparison

The **Decrypt - Replace Cert** decryption action is provided as an alternative to **Decrypt - Known Key** for these reasons:

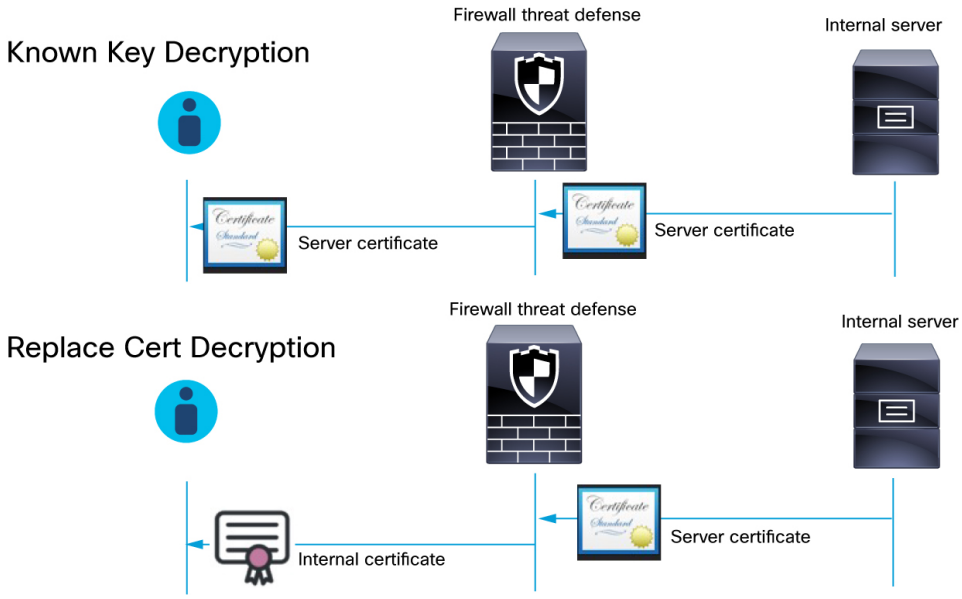
- Certificates can expire or need to be replaced anytime.

  If you choose the **Decrypt - Known Key** action, decryption can stop and the rule can stop functioning when a certificate is replaced for any reason.

- With **Decrypt - Replace Cert**, the connection matches and traffic is decrypted. Deploying the certificate on the device hours or days later is not an issue.

- **Decrypt - Replace Cert** gives you the flexibility to have separate internal and external facing certificates. For example, you could use `payroll.example.local` as an internal server hostname, but use `payroll.example.com` as an externally-facing hostname. You can also use an Enterprise PKI certificate for an internal server but use DigiCert or some other public provider for my external certificate.

The figure shows the difference between **Decrypt - Replace Cert** and **Decrypt - Known Key** decryption actions.



This table provides more information.

	Known Key	Replace Cert
Must match certificate installed on internal server or resource	Yes	No
The device must trust the internal certificate (that is, it must be chained to Trusted CAs list in the decryption policy)	No	No
Certificate hostnames can be different between device-installed certificate and that on the internal server	No, certificates must be identical. (Although multiple SANs can be on the certificates to accommodate different internal and external host names.)	Yes, certificates and certificate host names can be completely different, as long as the internal certificate is trusted by the user's device.
Works if the internal server certificate is expired	Yes – As long as it matches the certificate installed on the device.	No – an expired certificate causes broken trust, and therefore a mismatch.

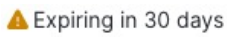
Replace an internal certificate (decrypt - replace cert only)

Replace expiring internal certificates in decryption policies to maintain secure encrypted traffic inspection capabilities.

This task discusses how to replace the internal certificate used in an incoming decryption rule using the Cloud-Delivered Firewall Management Center.

You can also replace the certificate using the API as discussed in *Cloud-Delivered Firewall Management Center REST API Quick Start Guide*.

The system indicates that an internal certificate is *expiring* when today's date is within 30 days of its expiration date. The following figure shows an example.



Before you begin

You must choose the **Decrypt - Replace Cert** decryption rule action in one of these ways:

- In a rule-based decryption policy by clicking the action from the rule's **Action** list.
- Create a standard decryption policy, which always uses the **Decrypt - Replace Cert** rule action.

Follow these steps to replace an internal certificate using the Decrypt - Replace Cert rule action:

Procedure

Step 1 Log in to Security Cloud Control if you haven't already done so.

Step 2 Click **Firewall**.

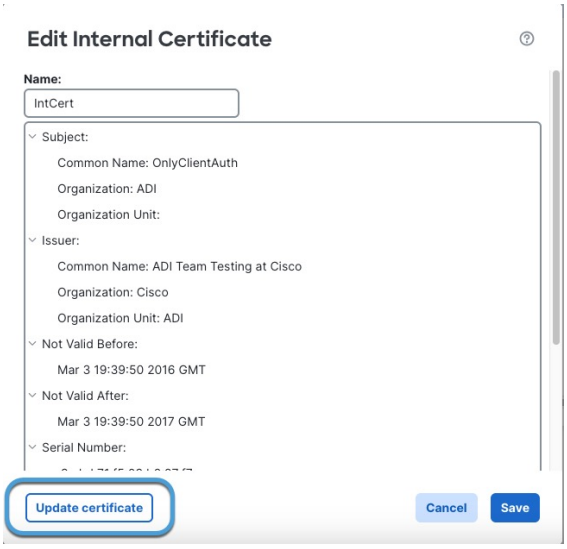
Step 3 Click **Administration** > **(name of** Cloud-Delivered Firewall Management Center**)** > **Objects** > **PKI** > **Internal Certs**.

Step 4 Locate the internal certificate that is expired or expiring (for example, an expired certificate is indicated by 

).

Step 5 Click **Edit** (✎).

Step 6 Click **Update Certificate** as this figure shows.



Step 7 Enter the requested information; see Upload an internal certificate for inbound protection, on page 2177 for more information.

Step 8 Follow the prompts on your screen to complete the action.

Decrypt - Resign rule action

The **Decrypt - Resign** rule action is a decryption rule action that enables you to decrypt *outgoing* traffic so the traffic can be subjected to deep inspection.

Select internal certificate objects

This task enables you to create a decryption rule with a **Decrypt - Known Key** action by selecting and assigning internal certificates to decrypt server traffic. You can add, manage, and save certificate selections for decryption rules in your configuration.

Use this procedure when configuring decryption rules that require internal certificates to decrypt traffic from specific servers.

Before you begin

To create a decryption rule with a **Decrypt - Known Key** action, you must add internal certificates for the servers to decrypt traffic. For more information, see Internal certificate objects, on page 2996.

Procedure

Step 1 Click **Edit** ( ) in the field next to the decryption rule **Action**.

Step 2 In the Select Internal Certificate Objects dialog box, do the following:
Available Certificates list: Click the name of a certificate and click **Add to Rule** to use it in this rule. You can add one or more certificates to the rule.

Selected Certificates list: Displays internal certificates used in the rule. To remove a certificate, click **Delete** ( ).

Step 3 To save these changes, click **OK**.

Monitor TLS/SSL hardware acceleration

TLS/SSL hardware acceleration monitoring is a set of topics that discuss how to monitor the status of TLS/SSL hardware acceleration.

Informational counters

Monitor TLS/SSL hardware acceleration performance and interpret KEY counter values, such as PRIV_KEY_RECV and SECU_PARAM_RECV, to assess system load and connection activity using the **show counters** command.

If a system under load is working well, you should see large counts for the following counters. Because there are two sides to the tracker process per connection, you can see these counters increase by two per connection. The PRIV_KEY_RECV and SECU_PARAM_RECV counters are the most important, and are highlighted. The CONTEXT_CREATED and CONTEXT_DESTROYED counters relate to the allocation of cryptographic chip memory.

```
> show counters
Protocol      Counter                          Value   Context
SSLENC        CONTEXT_CREATED                 258225   Summary
SSLENC        CONTEXT_DESTROYED               258225   Summary
TLS_TRK       OPEN_SERVER_SESSION             258225   Summary
TLS_TRK       OPEN_CLIENT_SESSION             258225   Summary
TLS_TRK       UPSTREAM_CLOSE                  516450   Summary
TLS_TRK       DOWNSTREAM_CLOSE                516450   Summary
TLS_TRK       FREE_SESSION                    516450   Summary
TLS_TRK       CACHE_FREE                      516450   Summary
TLS_TRK       PRIV_KEY_RECV                   258225   Summary
TLS_TRK       NO_KEY_ENABLE                   258225   Summary
TLS_TRK       SECU_PARAM_RECV                 516446   Summary
TLS_TRK       DECRYPTED_ALERT                 258222   Summary
TLS_TRK       DECRYPTED_APPLICATION         33568976   Summary
TLS_TRK       ALERT_RX_CNT                    258222   Summary
TLS_TRK       ALERT_RX_WARNING_ALERT          258222   Summary
TLS_TRK       ALERT_RX_CLOSE_NOTIFY           258222   Summary
TCP_PRX       OPEN_SESSION                    516450   Summary
TCP_PRX       FREE_SESSION                    516450   Summary
TCP_PRX       UPSTREAM_CLOSE                  516450   Summary
TCP_PRX       DOWNSTREAM_CLOSE                516450   Summary
TCP_PRX       FREE_CONN                       258222   Summary
TCP_PRX       SERVER_CLEAN_UP                 258222   Summary
TCP_PRX       CLIENT_CLEAN_UP                 258222   Summary
```

Alert counters

Alert counters identify potential issues based on standardized alert types.

These counters are based on the TLS 1.2 specification. FATAL or BAD alerts could indicate issues. However, ALERT_RX_CLOSE_NOTIFY is normal.

For details, see RFC 5246 section 7.2.

```
TLS_TRK       ALERT_RX_CNT                      311   Summary
TLS_TRK       ALERT_TX_CNT                        2   Summary
TLS_TRK       ALERT_TX_IN_HANDSHAKE_CNT           2   Summary
TLS_TRK       ALERT_RX_IN_HANDSHAKE_CNT           2   Summary
TLS_TRK       ALERT_RX_WARNING_ALERT            308   Summary
TLS_TRK       ALERT_RX_FATAL_ALERT                3   Summary
TLS_TRK       ALERT_TX_FATAL_ALERT                2   Summary
TLS_TRK       ALERT_RX_CLOSE_NOTIFY             308   Summary
TLS_TRK       ALERT_RX_BAD_RECORD_MAC             2   Summary
TLS_TRK       ALERT_TX_BAD_RECORD_MAC             2   Summary
TLS_TRK       ALERT_RX_BAD_CERTIFICATE            1   Summary
```

Error counters

Error counters help you interpret system performance data to identify potential issues and optimize performance using the show counters command.

Fatal counters

Monitor and interpret TLS/SSL hardware acceleration fatal counters to identify serious errors, assess system health, and troubleshoot persistent crypto acceleration issues on Cisco devices. Reset or review counters for proactive system maintenance and support.

Fatal counters indicate serious errors. These counters should be at or near 0 on a healthy system. This is an example list of the fatal counters.

```
> show counters
Protocol     Counter                          Value   Context
CRYPTO       RING_FULL                            1   Summary
CRYPTO       ACCELERATOR_CORE_TIMEOUT             1   Summary
CRYPTO       ACCELERATOR_RESET                    1   Summary
CRYPTO       RSA_PRIVATE_DECRYPT_FAILED           1   Summary
```

The RING_FULL counter is not a fatal counter, but indicates how often the system overloaded the cryptographic chip. The ACCELERATOR_RESET counter is the number of times the TLS crypto acceleration process failed unexpectedly, which also causes the failure of pending operations, which are the numbers you see in ACCELERATOR_CORE_TIMEOUT and RSA_PRIVATE_DECRYPT_FAILED.

If you have persistent problems, disable TLS crypto acceleration (or **config hwCrypto disable**) and work with Cisco TAC to resolve the issues.

Note You can do additional troubleshooting using the **show snort tls-offload** and **debug snort tls-offload** commands. Use the **clear snort tls-offload** command to reset the counters displayed in the **show snort tls-offload** command to zero.

Troubleshoot Rule-based decryption rules

This topic discusses how to troubleshoot rule-based decryption rules.

TLS/SSL oversubscription

TLS/SSL oversubscription is a state where a managed device is overloaded with TLS/SSL traffic. Any managed device can experience TLS/SSL oversubscription but only managed devices that support TLS crypto acceleration provide a configurable way to handle it.

TLS/SSL oversubscription details

When a managed device with TLS crypto acceleration enabled is oversubscribed, any packet received by the managed device is acted on according to the setting for **Handshake Errors** in the decryption policy's **Undecryptable Actions**:

- Inherit default action
- Do not decrypt
- Block
- Block with reset

If the setting for **Handshake Errors** in the decryption policy's **Undecryptable Actions** is **Do Not decrypt** and the associated access control policy is configured to inspect the traffic, inspection occurs; decryption does *not* occur.

Troubleshoot TLS/SSL oversubscription

Monitor and troubleshoot TLS/SSL oversubscription by configuring decryption policies, enabling logging, and adding SSL Flow Flags to connection event views in Cloud-Delivered Firewall Management Center.

If your managed device has TLS crypto acceleration enabled, you can view connection events to determine whether or not the devices are experiencing SSL oversubscription. You must add at least the **SSL Flow Flags** event to the table view of connection events.

Before you begin

- Configure a decryption policy with a setting for **Handshake Errors** on **Undecryptable Actions** page.

  For more information, see Set default handling for undecryptable traffic, on page 2181.

- Enable logging for your SSL rules as discussed in the section on logging decryptable connections in decryption rules in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.

Follow these steps to troubleshoot TLS/SSL oversubscription:

Procedure

Step 1 Click **Analysis** > **Unified Events**

Step 2 Click the colum picker ( ) icon and add additional columns **TLS Flow Flags** and **TLS Flow Messages**.

Step 3 Click **Apply**.
TLS/SSL oversubscription is indicated by the values of `ERROR_EVENT_TRIGGERED` and `OVER_SUBSCRIBED` in the **SSL Flow Flags** column.

Step 4 If TLS/SSL oversubscription is occurring, log in to the managed device and enter any of the following commands:

Command	Result
show counters	If the value of **TCP_PRX BYPASS_NOT_ENOUGH_MEM** is large, consider upgrading your device to one with a larger capacity for SSL traffic or use **Do NOT Decrypt** rules for lower priority encrypted traffic.
show snort TLS-offload	If the value of **BYPASS_NOT_ENOUGH_MEM** is large, consider upgrading your device to one with a larger capacity for SSL traffic or use **Do NOT Decrypt** rules for lower priority encrypted traffic.

TLS heartbeat

A TLS heartbeat extension enables applications to confirm active TLS or DTLS connections.

TLS heartbeat handling

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by RFC6520. TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

When a managed device with TLS crypto acceleration enabled encounters a packet that uses the TLS heartbeat extension, the managed device takes the action specified by the setting for **Decryption Errors** in the decryption policy's **Undecryptable Actions**:

- Block
- Block with reset

Related Topics

Troubleshoot TLS heartbeat, on page 2230

Troubleshoot TLS heartbeat

This task helps you identify TLS heartbeat activity in your network to determine if applications are using the TLS heartbeat extension.

If your managed device has TLS crypto acceleration enabled, you can view connection events to determine whether or not the devices are seeing traffic with the TLS heartbeat extension. You must add at least the **SSL Flow Messages** event to the table view of connection events.

TLS heartbeat is indicated by the value of `HEARTBEAT` in the **SSL Flow Messages** column in the table view of connection events.

Before you begin

To determine if applications in your network use TLS heartbeat, first perform these tasks:

- Configure an decryption policy with a setting for **Decryption Errors** on **Undecryptable Actions** page.

  For more information, refer to Set default handling for undecryptable traffic, on page 2181.

- Enable logging for your SSL rules as discussed in Secure Firewall Management Center and Threat Defense Management Network Administration.

Follow these steps to troubleshoot TLS heartbeat:

Procedure

Step 1 Click **Analysis** > **Unified Events**

Step 2 Click the colum picker ( ) icon and add additional columns **TLS Flow Flags** and **TLS Flow Messages**.

Step 3 Click **Apply**.
TLS heartbeat is indicated by the value of `HEARTBEAT` in the **SSL Flow Messages** column.

Step 4 If applications in your network use TLS heartbeat, refer to Rule-based decryption rules guidelines and limitations, on page 2188.

TLS/SSL pinning

TLS/SSL pinning is a security mechanism that

- validates server certificates against a predetermined set of trusted certificates or keys
- prevents man-in-the-middle attacks by rejecting connections to servers with unexpected certificates, and
- requires applications to maintain a local store of trusted certificate information for validation.

TLS/SSL pinning implementation details

Some applications use a technique referred to as *TLS/SSL pinning* or *certificate pinning*, which embeds the fingerprint of the original server certificate in the application itself. As a result, if you configured a decryption rule with a **Decrypt - Resign** action, when the application receives a resigned certificate from a managed device, validation fails and the connection is aborted.

To confirm that TLS/SSL pinning is occurring, attempt to log in to a mobile application like Facebook. If a network connection error is displayed, log in using a web browser. (For example, you *cannot* log in to a Facebook mobile application but *can* log in to Facebook using Safari or Chrome.) You can use Firepower Management Center connection events as further proof of TLS/SSL pinning

Note TLS/SSL pinning is not limited to mobile applications.

If applications in your network use SSL pinning, see TLS/SSL certificate pinning guidelines, on page 2194.

Related Topics

Troubleshoot TLS/SSL pinning, on page 2231

Troubleshoot TLS/SSL pinning

This task helps you identify and confirm whether applications in your network are using TLS/SSL pinning by analyzing SSL connection events and flow DATA.

You can view connection events to determine whether or not the devices are experiencing TLS/SSL pinning. You must add at least the **SSL Flow Flags** and **SSL Flow Messages** columns to the table view of connection events.

Before you begin

- Enable logging for your decryption rules as discussed in the section on logging decryptable connections in decryption rules in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.
- Log in to a mobile application like Facebook. If a network connection error displays, log in to Facebook using Chrome or Safari. If you *can* log in using a web browser but not the native application, TLS/SSL pinning is likely occurring.

Follow these steps to troubleshoot TLS/SSL pinning:

Procedure

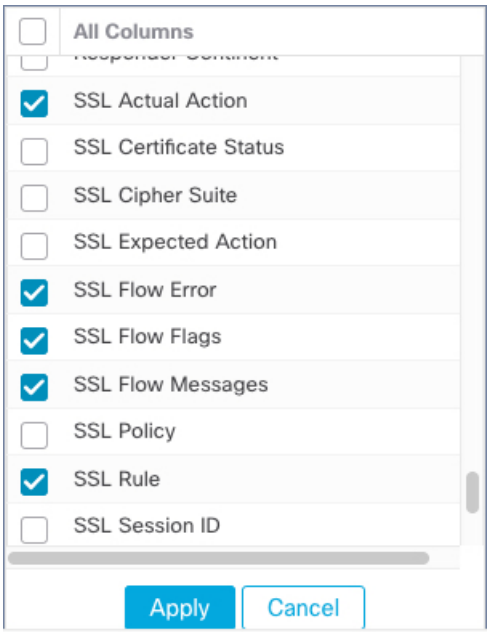
Step 1 Click **Analysis** > **Connections** > **Events**.

Step 2 Click **Table View of Connection Events**.

Step 3 Click **x** on any column in the connection events table to add additional columns for at least **SSL Flow Flags** and **SSL Flow Messages**.



The following example shows adding the **SSL Actual Action**, **SSL Flow Error**, **SSL Flow Flags**, **SSL Flow Messages**, **SSL Policy**, and **SSL Rule** columns to the table of connection events.



The columns are added in the order discussed in the section on connection and security intelligence event fields in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.

Step 4 Click **Apply**.

Step 5 Analyze the SSL flow DATA to identify SSL pinning behavior.

Step 6 If you determine that applications in your network use SSL pinning, see Rule-based decryption rules guidelines and limitations, on page 2188.

What to do next

You can use TLS/SSL connection events to confirm TLS/SSL pinning is occurring by looking for any of the following:

- Applications that send an SSL ALERT Message as soon as the CLIENT receives the SERVER_HELLO, SERVER_CERTIFICATE, SERVER_HELLO_DONE message from the SERVER, followed by a TCP Reset, exhibit the following symptoms. (The ALERT, `Unknown CA (48)`, can be viewed using a packet capture.)
  - The SSL Flow Flags column displays `ALERT_SEEN` but *not* `APP_DATA_C2S` or `APP_DATA_S2C`.
  - If your managed device has SSL hardware acceleration enabled, the SSL Flow Messages column typically displays: `CLIENT_ALERT`, `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`.
  - If your managed device doesn't support SSL hardware acceleration or if the feature is disabled, the SSL Flow Messages column typically displays: `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`.
  - `Success` is displayed in the SSL Flow Error column.
- Applications that send no alerts but instead send TCP Reset after the SSL handshake is FINISHED exhibit the following symptoms:
  - The SSL Flow Flags column does *not* display `ALERT_SEEN`, `APP_DATA_C2S`, or `APP_DATA_S2C`.
  - If your managed device has SSL hardware acceleration enabled, the SSL Flow Messages column typically displays: `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`, `CLIENT_KEY_EXCHANGE`, `CLIENT_CHANGE_CIPHER_SPEC`, `CLIENT_FINISHED`, `SERVER_CHANGE_CIPHER_SPEC`, `SERVER_FINISHED`.
  - If your managed device doesn't support SSL hardware acceleration or if the feature is disabled, the SSL Flow Messages column typically displays: `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`, `CLIENT_KEY_EXCHANGE`, `CLIENT_CHANGE_CIPHER_SPEC`, `CLIENT_FINISHED`,`SERVER_CHANGE_CIPHER_SPEC`, `SERVER_FINISHED`.
  - `Success` is displayed in the SSL Flow Error column.

Related Topics

Troubleshoot unknown or bad certificates or certificate authorities, on page 2233

Troubleshoot unknown or bad certificates or certificate authorities

Enable logging for TLS/SSL decryption rules to monitor connection events and identify certificate-related issues. This helps you diagnose problems with certificate authorities, bad certificates, or unknown certificates in your TLS/SSL connections.

You can view connection events to determine whether or not the devices are experiencing unknown certificate authorities, bad certificates, or unknown certificates. This procedure can also be used if a TLS/SSL certificate has been pinned. You must add at least the **SSL Flow Flags** and **SSL Flow Messages** columns to the table view of connection events.

Before you begin

- Set up a decryption rule.

- Enable logging for your decryption rules as discussed in the section on logging decryptable connections in decryption rules in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.

Follow these steps to troubleshoot unknown or bad certificates or certificate authorities:

Procedure

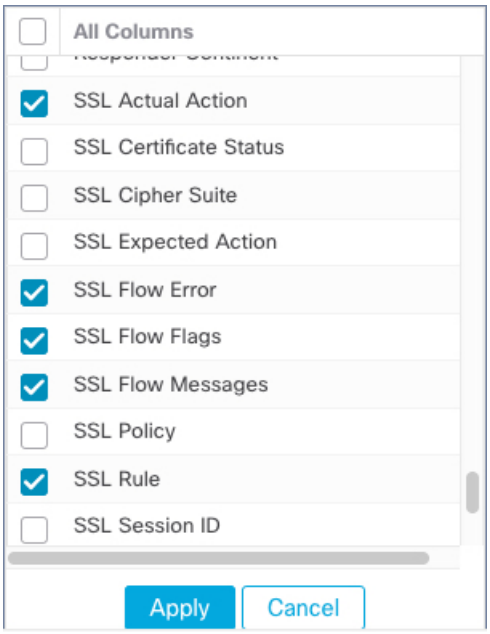
Step 1 Click **Analysis** > **Connections** > **Events**.

Step 2 Click **Table View of Connection Events**.

Step 3 Click **x** on any column in the connection events table to add additional columns for at least **SSL Flow Flags** and **SSL Flow Messages**.



The following example shows adding the **SSL Actual Action**, **SSL Flow Error**, **SSL Flow Flags**, **SSL Flow Messages**, **SSL Policy**, and **SSL Rule** columns to the table of connection events.



The columns are added in the order discussed in the section on connection and security intelligence event fields in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.

Step 4 Click **Apply**.

Step 5 Review the SSL flow flags to determine if a certificate or certificate authority is bad or missing.

SSL flow flag	Meaning
CLIENT_ALERT_SEEN_UNKNOWN_CA	Indicates a valid certificate chain or partial chain was received by an SSL client application, but the certificate was not accepted because the CA certificate could not be located or could not be matched with a known, trusted CA. This message always indicates an unrecoverable error.
CLIENT_ALERT_SEEN_BAD_CERTIFICATE	A certificate was corrupt, contained signatures that did not verify correctly, or had other problems.
CLIENT_ALERT_SEEN_CERTIFICATE_UNKNOWN	Some other (unspecified) issue arose in processing the certificate, rendering it unacceptable.

Verify TLS/SSL cipher suites

This task helps you resolve cipher suite incompatibility errors that occur when saving a decryption rule that has cipher suite conditions. The error indicates that one or more of the cipher suites you chose for the decryption rule condition are incompatible with the certificate used in the decryption rule.

This topic discusses actions you must take if you see the following error when saving a decryption rule that has cipher suite conditions:

```
Traffic cannot match this rule; none of your selected cipher suites contain a signature
algorithm that matches the resigning CA's signature algorithm
```

Before you begin

To resolve the issue, you must have access to the certificate you're using.

Note The tasks in this topic assume knowledge of how TLS/SSL encryption works.

Follow these steps to verify TLS/SSL cipher suites:

Procedure

Step 1 When you attempt to save a decryption rule rule with either **Decrypt - Resign** or **Decrypt - Known Key** with specified cipher suites, the following error is displayed:

Example:

```
Traffic cannot match this rule; none of your selected cipher suites contain a
signature algorithm that the resigning CA's signature algorithm
```

Step 2 Locate the certificate you're using to decrypt traffic and, if necessary, copy the certificate to a system that can run openssl commands.

Step 3 Run the following command to display the signature algorithm used by the certificate:

Example:

openssl x509 -in *CertificateName* **-text -noout**

The first few lines of output are displayed similar to the following:

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 4105 (0x1009)
    Signature Algorithm: ecdsa-with-SHA256
```

Step 4 The **Signature algorithm** tells you the following:

- The cryptographic function used (in the preceding example, **ECDSA** means Elliptic Curve Digital Signature Algorithm).
- The hash function used to create a digest of the encrypted message (in the preceding example, **SHA256**).

Step 5 Search a resource such as *OpenSSL at University of Utah* for cipher suites that match those values. The cipher suite must be in RFC format.
You can also search a variety of other sites, such as Server Side TLS at the Mozilla wiki or Appendix C of RFC 5246. Cipher Suites in TLS/SSL (Schannel SSP) in Microsoft documentation has a detailed explanation of cipher suites.

Step 6 If necessary, translate the OpenSSL name to an RFC name that the Firepower Management System uses. See the RFC mapping list on the on the `https://testssl.sh` site.

Step 7 The previous example, **ECDSA-with-SHA256**, can be found in the Modern Compatibility List on the Mozilla wiki.

Choose only cipher suites that have **ECDSA** and **SHA-256** in the name. These cipher suites follow:

```
ECDHE-ECDSA-AES128-GCM-SHA256
ECDHE-ECDSA-AES128-SHA256
```

Find the corresponding RFC cipher suite on RFC mapping list. These cipher suites follow:

```
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
```

Step 8 Add the preceding cipher suites to your decryption rule.

Troubleshoot using crypto archives

Crypto issues are difficult to triage. Crypto archives help you to troubleshoot these issues. A crypto archive contains crypto session information about the crypto request, peer information, the component that sent the crypto request, and the timed-out crypto session information.

About crypto archives

Threat defense does not save keys and initialization vectors (IVs) for the session. For SSL and IPsec, you can also view these information:

- For SSL: Session SSL version, source, destination IP addresses, and ports.
- For IPsec: IPsec security association information.

A ring can hold 2000 crypto command entries. Threat defense pushes the crypto command in one of the rings and pulls out the result after completing the crypto request. Crypto archive files now have the timed-out crypto request's ring and entry index. The ring and its entry index help in troubleshooting problematic crypto commands.

The crypto archive has two formats: a text file and a binary file. You can use the **debug menu ctm 103** command to decode the binary file. The crypto archive files are available at disk0:/crypto_archive.

For example:

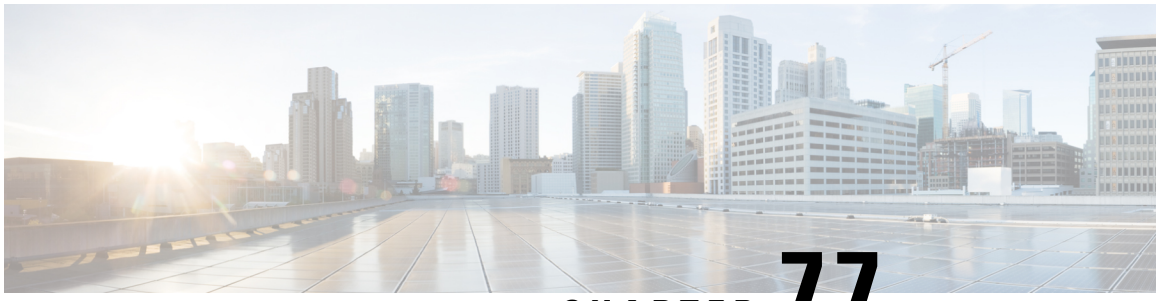
```
FTD# debug menu ctm 103 crypto_eng0_arch_4.bin
[Nitrox V Archive Header v1.0 Info]
ASA Image Version: PIX (9.20) #0: Tue Mar 29 16:20:30 GMT 2022
...
SE SSL microcode: CNN5x-MC-SE-SSL-0011
AE microcode: CNN5x-MC-AE-MAIN-0002
Crypto Engine 0
Crash type: SE Ring Timeout
...
Core Soft Resets: 11
...
Timeout Ring (SE): 12
Timeout Entry: 642
SE TIMEOUT:
Core SE 6 Touts: 2
Core SE 8 Touts: 2
Core SE 12 Touts: 4
Core SE 32 Touts: 2
Core SE 37 Touts: 1
.....
[Timeout Session Info]
Active: TRUE
Sync: FALSE
Callback: TRUE
Saved Callback: FALSE
Commands in progress: 1
Engine : hardware
Device : n5 (Nitrox V)
Session : ssl
Priority: normal
NP VPN context handle : 0x00000000
Flag : 0
vcid : 0
Block size : 2050
async cb ring index: 0
tls offload rsa: FALSE
Session context:
SSL Version : dtls1.2
SSL Context Type : handshake
Encryption Mode : gcm
Auth Algorithm : null
Hash Algorithm : none
Key Size : 32
SSL V : dtls1.2
Source IP : 82.1.2.2
Source Port : 51915
Dest IP : 82.29.155.32
Dest Port : 443
```

In this example, the highlighted information shows the timeout ring, the crash time (timeout entry), and SSL session information.

Supported devices for crypto archives

These devices with Nitrox V crypto accelerator support crypto archives:

- Cisco Firepower 3105, 3110, 3120, 3130, 3140
- Cisco Firepower 4112, 4115, 4125, 4145
- Cisco Firepower 9300 SM-40, SM-48 and SM-56
- Cisco Secure Firewall 4200



CHAPTER 77

Encrypted Visibility Engine

Encrypted Visibility Engine (EVE) is used to identify client applications and processes utilizing Transport Layer Security (TLS) encryption. It enables visibility and allows administrators to take actions and enforce policy within their environments. The EVE technology can also be used to identify and stop malware.

- Encrypted Visibility Engine, on page 2239
- How EVE works, on page 2240
- Indications of compromise events, on page 2241
- QUIC fingerprinting in EVE, on page 2241
- Configure EVE, on page 2242
- Configure EVE exception rules, on page 2243
- Use case - Block Traffic Based on the EVE Threat Confidence Score, on page 2247

Encrypted Visibility Engine

The encrypted visibility engine (EVE) is used to provide more visibility into the encrypted sessions without the need to decrypt them. These insights into encrypted sessions are obtained by Cisco's open-source library that is packaged in Cisco's vulnerability database (VDB). The library fingerprints and analyzes incoming encrypted sessions and matches it against a set of known fingerprints. This database of known fingerprints is also available in the Cisco VDB.

Key capabilities

Important features of Encrypted Visibility Engine (EVE) include the following:

- Access control policy actions on traffic using information derived from EVE
- Vulnerability Database (VDB) integration with Cisco Secure Firewall for assigning applications to EVE-detected processes with high confidence values
- Custom application detector creation for mapping EVE-detected processes to user-defined applications and overriding built-in process confidence values
- Detection of the operating system type and version of clients that create Client Hello packets in encrypted traffic
- Quick UDP Internet Connections (QUIC) traffic fingerprinting and analysis with server name display in the URL field of the **Connection Events** page

For custom application detector configuration, see the **Configuring Custom Application Detectors** and **Specifying EVE Process Assignments** sections in the **Application Detection** chapter of the Cisco Secure Firewall Management Center Device Configuration Guide.

Note The encrypted visibility engine feature is supported only on Cloud-Delivered Firewall Management Center-managed devices running Snort 3. This feature is not supported on Snort 2 devices and Firewall Device Manager-managed devices.

Attention To use EVE on Cloud-Delivered Firewall Management Center, you must have a valid IPS license on your device. In the absence of a IPS license, the policy displays a warning and deployment is not allowed.

Note

- EVE can detect the operating system type and version of SSL sessions. Normal usage of the operating system, such as running applications and package management software, can trigger OS detection. To view client OS detection, in addition to enabling the EVE toggle button, you must enable **Hosts** under **Policies** > **Network Discovery**. To view a list of possible operating systems on the host IP address, click **Analysis** > **Hosts heading** > **Network Map**, and then choose the required host.
- After enabling EVE for your access control policy, ensure that you have turned on logging for the access control rules within that policy to display the expected results on the EVE dashboard whenever any specific rule conditions are met. For more information on how to turn on logging, see Create and Edit Access Control Rules.

How EVE works

Encrypted Visibility Engine (EVE) provides the ability to identify and control applications without enabling TLS decryption. By using fingerprints of known malicious processes, EVE technology can also be used to identify and block encrypted malicious traffic without outbound decryption.

Summary

The key components involved in how EVE works are:

- Client Hello inspection: EVE examines the initial TLS handshake packet to create client fingerprints
- Fingerprint database: Contains over 5,000 identified client processes mapped to applications
- Machine learning technology: Processes over one billion TLS fingerprints and over 10,000 malware samples daily
- Cisco Vulnerability Database (VDB): Delivers updated fingerprints to customers

Workflow

These stages describe how EVE identifies and processes client applications:

1. EVE inspects the Client Hello portion of the TLS handshake to identify client processes. The Client Hello is the initial data packet that is sent to the server. This gives a good indication of the client process on the host.
2. The system combines the fingerprint with other data such as destination IP address to provide the basis for EVE's application identification. By identifying specific application fingerprints during the TLS session establishment, the system can identify the client process and take appropriate action, such as allowing or blocking it.
3. If EVE does not recognize a fingerprint, it identifies the client application and estimates the threat score of the first flow using the destination details, such as IP address, port, and server name. At this point, the status of the fingerprints are randomized and the status can be viewed in the debug logs. For subsequent flows with the same fingerprint, EVE skips reanalysis and marks the fingerprint status as unlabeled. If you intend to block traffic based on EVE's Low or Very Low score thresholds, the initial flow is blocked. Future flows are allowed once the application's fingerprint is cached.
4. Through machine learning (ML) technology, Cisco processes over one billion TLS fingerprints and over 10000 malware samples daily to create and update EVE fingerprints. These updates are then delivered to customers using Cisco Vulnerability Database (VDB) package.

Indications of compromise events

An Indication of Compromise (IoC) event is a security detection mechanism that

- identifies connection events with a very high malware confidence level, as reported by EVE,
- triggers for encrypted sessions generated from a host using a malicious client, and
- provides information such as the IP address, MAC address, operating system information of the malicious host, and timestamp of the suspicious activity.

A session with an Encrypted Visibility Threat Confidence score of 'Very High' as seen in connection events generates an IoC event. You must enable **Hosts** from **Policies** > **Network Discovery**. In the Cloud-Delivered Firewall Management Center, view the IoC event existence from here:

- Click **Analysis** > **Unified Events** and choose the Encrypted Visibility fields and IoC field from the column picker option.

QUIC fingerprinting in EVE

QUIC fingerprinting in EVE is a Snort identification mechanism that

- detects applications over QUIC without enabling decryption,
- identifies malware without enabling decryption, and
- detects service applications for access control rule assignment based on the service detected over the QUIC protocol.

Configure EVE

Configure the Encrypted Visibility Engine (EVE) to detect client applications and monitor or block encrypted traffic based on threat confidence levels.

The Encrypted Visibility Engine (EVE) provides visibility into encrypted traffic by analyzing client application behavior and threat indicators. You can operate EVE in Monitor mode for detection only or Protect mode to actively block threats.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** ( ) next to the access control policy you want to edit.

Step 3 Choose **Encrypted Visibility Engine** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 On the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button.

Step 5 Choose the **Monitor** mode or the **Protect** mode.

- Choose the **Monitor** mode to detect client applications and monitor encrypted traffic.
- Choose the **Protect** mode to monitor and block encrypted traffic based on the threat confidence level of the client processes. You can use this mode to monitor and block malicious connections at two threat confidence levels:
  - **High**: Use this level to block connections with threat confidence levels ranging from High to Very High.
  - **Very High**: Use this level to block connections with threat confidence levels that are categorized as Very High.

Step 6 Click **Save** and then deploy the access control policy.

Note
To manage exceptions, refer to Configure EVE exception rules, on page 2243.

What to do next

Deploy configuration changes.

View Encrypted Visibility Engine events

This task allows you to access and view connection events generated by the Encrypted Visibility Engine (EVE) in the Cloud-Delivered Firewall Management Center.

After enabling the **Encrypted Visibility Engine** and deploying your access control policy, you can start sending live traffic through your system. You can view the logged connection events in the **Unified Events** page.

Perform this procedure to access the connection events in the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Click **Analysis** > **Unified Events**.

You can also view the connection events in the **Connection Events** page. Click **Analysis** > **Connections** > **Events** to access the **Connection Events** page.

The Encrypted Visibility Engine can identify the client process that initiated a connection and the operating system in the client, and indicate if the process contains malware or not.

Step 2 In the **Unified Events** page, explicitly enable these columns that are added for the Encrypted Visibility Engine:

- **Encrypted Visibility Process Name**
- **Encrypted Visibility Process Confidence Score**
- **Encrypted Visibility Threat Confidence**
- **Encrypted Visibility Threat Confidence Score**
- **Detection Type**
- **EVE Fingerprint**
- **EVE Process Name**
- **EVE Process Confidence Score**
- **EVE Threat Confidence**
- **EVE Threat Confidence Score**
- **Detection Type**

For information about these fields, see Connection and Security-Related Connection Event Fields in the *Cisco Secure Firewall Management Center Administration Guide*.

Note
On the **Connection Events** page, if processes are assigned applications, the **Detection Type** column displays **Encrypted Visibility Engine**, indicating that the client application was identified by the Encrypted Visibility Engine. Without application assignments to process names, the **Detection Type** column displays **AppID**, indicating that the engine that identified the client application was AppID.

Configure EVE exception rules

Configure Encrypted Visibility Engine (EVE) exception rules to ensure the continuity of trusted connections and services by bypassing EVE's block action. This allows trusted networks and processes to be exempted from EVE's block verdict based on the threat confidence level.

You can create an Encrypted Visibility Engine (EVE) exception rule to ensure the continuity of trusted connections and services by bypassing the EVE's block action. You can add attributes such as process names and destination IP address to the exception rule. For example, you may want to bypass EVE's block verdict for trusted networks. All the connections in the bypassed networks are exempted from EVE's block verdict based on the threat confidence level.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** ( ) next to the access control policy you want to edit.

Step 3 Choose **Encrypted Visibility Engine** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 On the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button.

Step 5 Choose the **Protect** mode to monitor and block encrypted traffic based on the threat confidence level of the client processes.

You can use this mode to monitor and block malicious connections at two threat confidence levels:

- **High**: Use this level to block connections with threat confidence levels ranging from High to Very High.
- **Very High**: Use this level to block connections with threat confidence levels that are categorized as Very High.

Step 6 Click **Manage exceptions** to view and add exception rules.

Step 7 On the **Encrypted Visibility Engine (EVE) Exception List** window, click **+Add Exception Rules** and add the required attributes.

a) Under the **Process Name** tab, enter an EVE-identified process name, and click **+Add** on the right side of the window.

You can add multiple process names to the same exception rule. The EVE exception list based on process names works only with EVE-identified process names, which are case and space sensitive.

b) Under the **Network Objects** tab, perform one of the following:

- Choose one or more network objects from the **Available Networks** list and add the same to the **Selected Source Network** or **Selected Destination Network** list.
- To create a new network object, click **+Create Network Object**.
  1. Enter a **Name** and an optional **Description**.
  2. Choose the required network type—**Host, Range, Network, or FQDN**. Enter the relevant IP address if you choose **Host, Range, or Network**. If you choose **FQDN**, enter the fully Qualified Domain Name(**FQDN**) and choose the required option from the **Lookup** drop-down list.
  3. If you want to allow configuration overrides, check the **Allow overrides** checkbox.
  4. Click **Add**.

c) To create a new dynamic attribute, click **+Create Dynamic Attribute.**

1. Enter a **Name** and an optional **Description**.

2. Click **Add**. You can configure this object using Cisco Secure Dynamic Attribute Connector (CSDAC) or Management Center APIs.

d) (Optional) In the **Comment** field available on all the tabs, you can enter a reason for adding the required network objects and dynamic attributes to the EVE exception rule.

Step 8 Click **Save** and then deploy the access control policy.



Note When a connection matches an exception rule, it bypasses the EVE's block verdict. You can view EVE's action in the **Connection Events** or **Unified Events** page. The **Reason** column header displays **EVE Exempted** for identification of such EVE-bypassed traffic.

Add exception rule from unified events

Add exception rules for connections that are blocked by EVE using the information in the Unified Events page.

Use the **Unified Events** page to add exception rules for connections that are blocked by EVE. The Cloud-Delivered Firewall Management Center adds an exception rule to the **Encrypted Visibility Engine (EVE) exception list** object. Note that the exception rules added to this list are applicable for all the access control policies that have EVE enabled.

Before you begin

Exception list is supported only from threat defense version 7.6.0 or later.

Follow these steps to add exception rules for connections that are blocked by EVE:

Procedure

Step 1 Click **Analysis** > **Unified Events**.

Step 2 In the **Reason** column with **Encrypted Visibility Block** as the reason, click the **Ellipsis**(⋮) icon inside the cell.

Step 3 Choose **Add EVE Exception Rule** from the drop-down list.

Step 4 In the **Encrypted Visibility Engine** window that is displayed, the rule is automatically added to the bottom of the exception list. You can review and make changes to the added rule before saving and deploying the configuration.

Upgrade EVE exception rules

On Secure Firewall version 7.7 and earlier, Encrypted Visibility Engine (EVE) exception rules are configured for each policy separately. From Secure Firewall version 10.0.0, the EVE exception list is part of the global domain. As a result, the EVE exception rules are configured in the global domain and applied to all the policies on which EVE is enabled, to block traffic.

EVE exception rule upgrade behavior

When you are upgrading the Management Center from version 7.7 to 10.0.0, all the EVE exception rules from the leaf domains that contain leaf domain network objects, are identified and stored. After the upgrade is complete:

- All EVE exception rules from global domain policies, as well as rules from leaf domain policies that reference global domain objects or inline IP addresses, are consolidated into a single global EVE exception list. Some policies may now include EVE exception rules that were not present before the upgrade as a result.
- All policies that contain EVE exception rules are marked as out-of-date.

If the exception rules from leaf domains contain leaf domain network or dynamic objects, these rules are removed during the upgrade process. The upgrade script log file records all the merged and deleted exception rules and includes the corresponding access control policy and domain for each rule. The log file is located at */var/log/sf/Cisco_Secure_FW_Mgmt_Center_Upgrade10.0.0/800_post/1114_eve_rules.pl.log*.

When you deploy the configuration for the first time after the upgrade, a warning message on the Management Center lists all the deleted EVE exception rules. The warning message also states that there could be possible traffic impact if the rules are not reconfigured in the global EVE exception list. The warning message appears only when you deploy the configuration for the first time after the upgrade is complete.

For Secure Firewall devices running version 7.7 and earlier that are mapped to Management Center running version 10.0.0, only **Very High** threat confidence connection events are sent to the **Security-Related Connection Events** table. For Secure Firewall devices running version 10.0.0, EVE **Blocked** and **Medium+** EVE threat confidence connection events are sent to the **Security-Related Connection Events** table.

Change management support during EVE upgrade:

- When you upgrade the Management Center to version 10.0.0, all active change management tickets that contain access control policies on which EVE is enabled will have their EVE exception rules automatically merged with the global EVE exception list.
- The merging of EVE exception rules with the global EVE exception list occurs regardless of the ticket's approval state. This ensures that no exception rules are lost during the upgrade.

EVE ticket preview generation behavior:

- If a change management ticket contains a policy that is locked and it contains only EVE-related modifications, such as EVE settings or exception rules, the EVE ticket preview will not be automatically regenerated after the upgrade.
- If the ticket contains other policy modifications in addition to EVE-related modifications, the EVE ticket preview will be generated normally.

Use case - Block Traffic Based on the EVE Threat Confidence Score

About Encrypted Visibility Engine

You can use the Encrypted Visibility Engine (EVE) to identify client applications and processes using Transport Layer Security (TLS) encryption. EVE provides more visibility into the encrypted sessions without decryption. Based on EVE's findings, administrators can enforce policy actions on the traffic within their environments. You can also use the EVE to identify and stop malware.

Benefits

Administrators can leverage and adjust EVE's threat score to block malicious encrypted traffic. If the probability that the incoming traffic is malicious, then based on the threat score, you can configure EVE to block the connection.

Sample Business Scenario

A large corporate network uses Snort 3 as its primary intrusion detection and prevention system. In a rapidly evolving threat landscape, adoption of robust network security measures is necessary and important. The security team uses EVE to enhance encrypted traffic inspection without the need to implement full man-in-the-middle (MITM) decryption. The EVE technology uses fingerprints of known malicious processes to identify and stop malware. Network administrators must have the flexibility to configure EVE's block traffic thresholds to block potentially malicious connections, which are based on their configured block thresholds.

Prerequisites

- You must be running management center 7.4.0 or later, and the managed threat defense must also be 7.4.0 or later.
- Ensure that you have a valid Intrusion Prevention System (IPS) license and Snort 3 is the detection engine.

High-Level Workflow

1. EVE analyzes the incoming traffic and gives a verdict on the probability of incoming traffic being malware or not.
2. If EVE detects incoming traffic to be malware with a certain level of confidence, you can configure EVE to block that traffic.
3. The packets are first checked for malware probability or threat score, and the threat score is compared with the block threshold that you have set.
4. If the threat score is higher than the configured threshold, EVE blocks the traffic.

5. If the threat score is lesser than the configured threshold, EVE takes no action.

Configure Block Thresholds in EVE

This procedure shows how to block potentially malicious traffic, based on the EVE threat confidence score of 90 percent or higher.

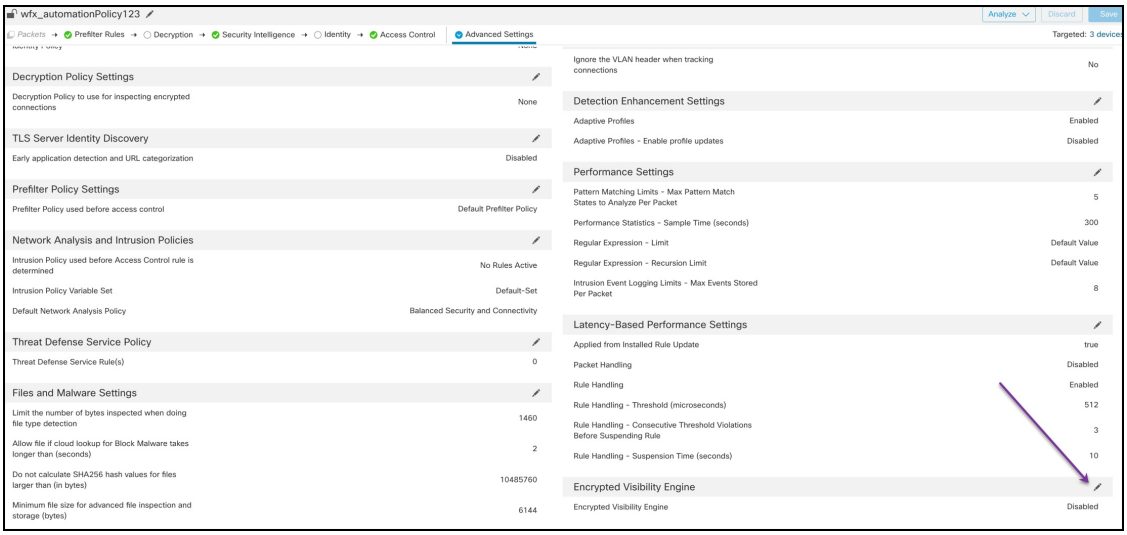
Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

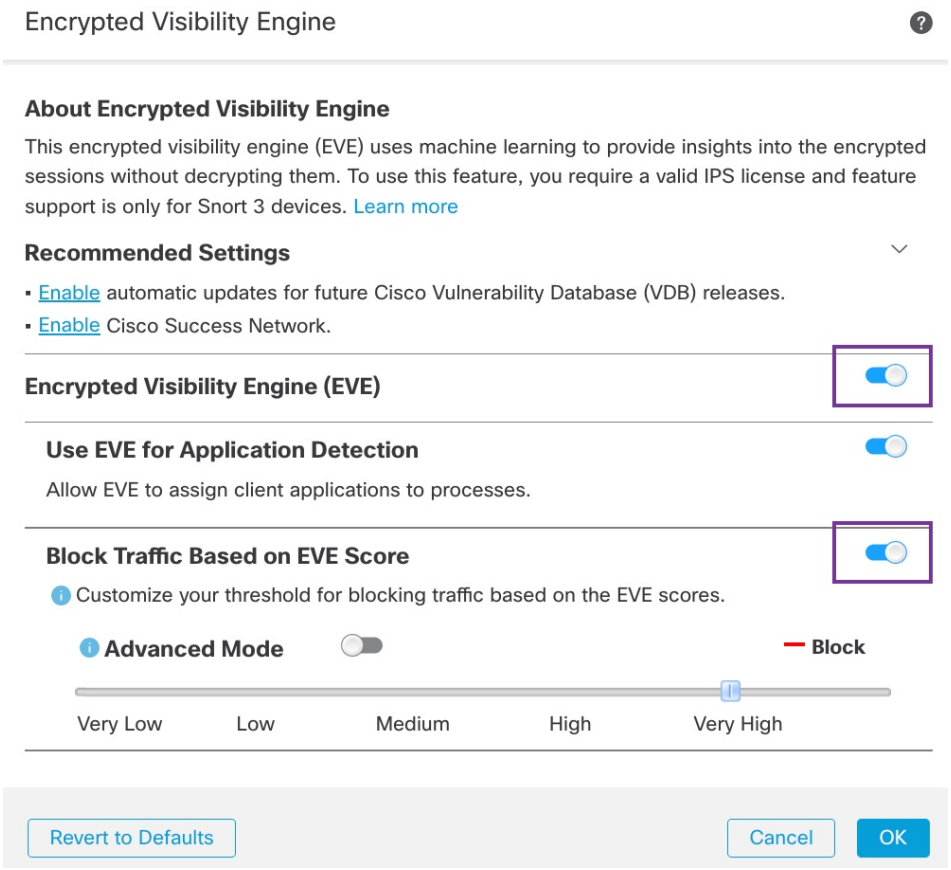
Step 3 Choose **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 Click **Edit** (✎) next to **Encrypted Visibility Engine**.



Step 5 In the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button.

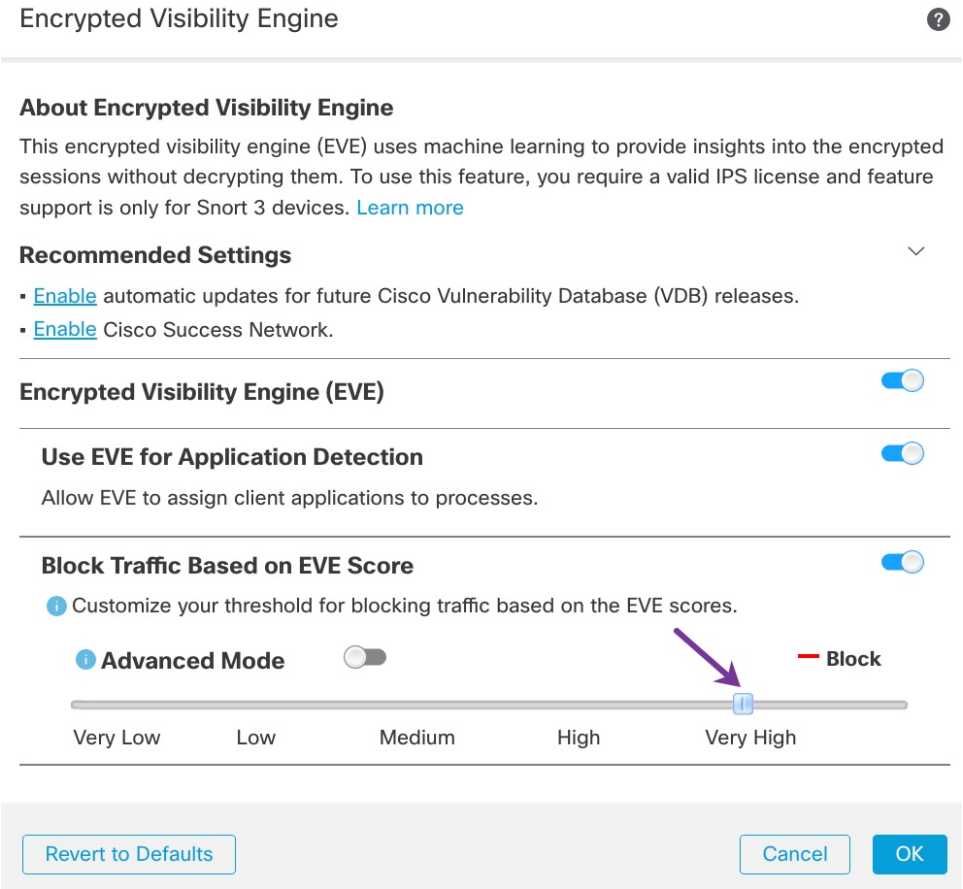
Step 6 Enable the **Block Traffic Based on EVE Score** toggle button. Any incoming traffic that is a potential threat is blocked by default.



Note
By default, the threshold at which malware is blocked is 99 percent, which means:

- If EVE detects the traffic to be malware with 99 percent confidence or higher, EVE blocks the traffic.
- If EVE detects the traffic to be malware with less than 99 percent confidence, EVE takes no action.

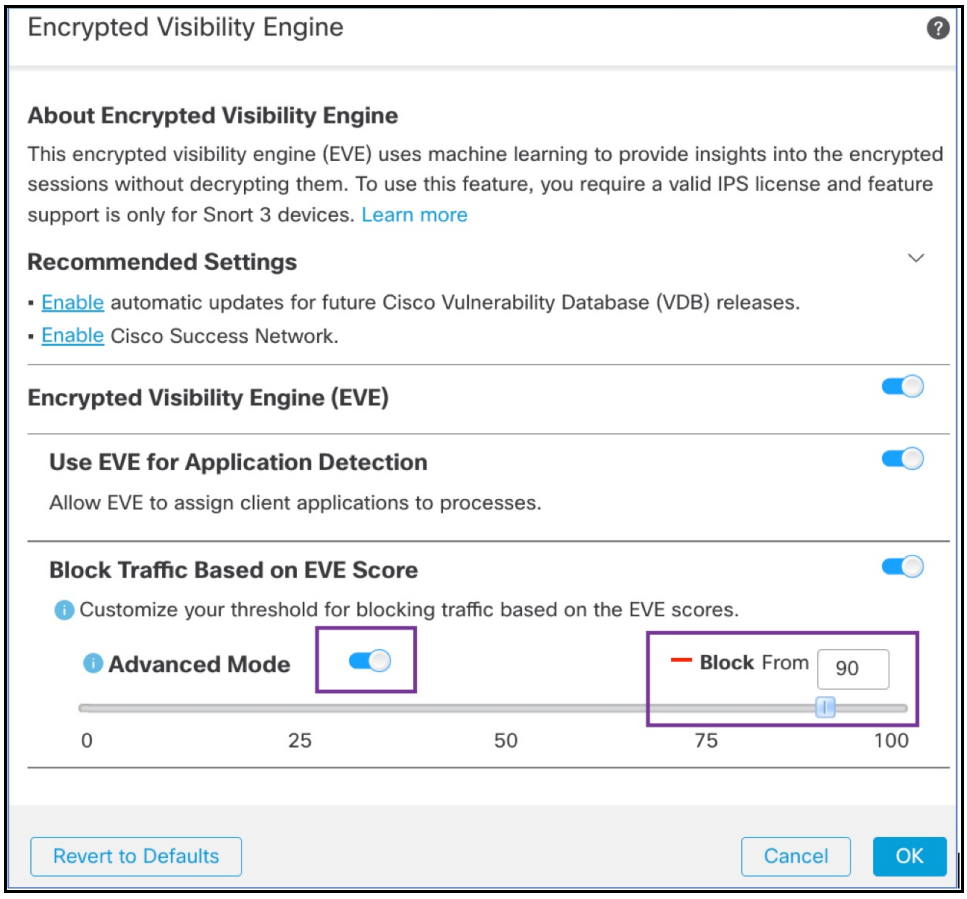
Step 7 Use the slider to adjust the threshold for blocking based on EVE threat confidence. This ranges from **Very Low** to **Very High**. In this example, the slider is set to **Very High**.



Step 8 For further granular control, enable the **Advanced Mode** toggle button. Now, you can assign a specific EVE Threat Confidence Score for blocking traffic. The default threshold is 99 percent.

Step 9 In this example, change the block threshold to **90** percent.

Attention
As a best practice, we recommend that you do not set the block threshold to below 50 percent to ensure optimum performance.



Step 10 Click **OK**.

Step 11 Click **Save**.

What to do next

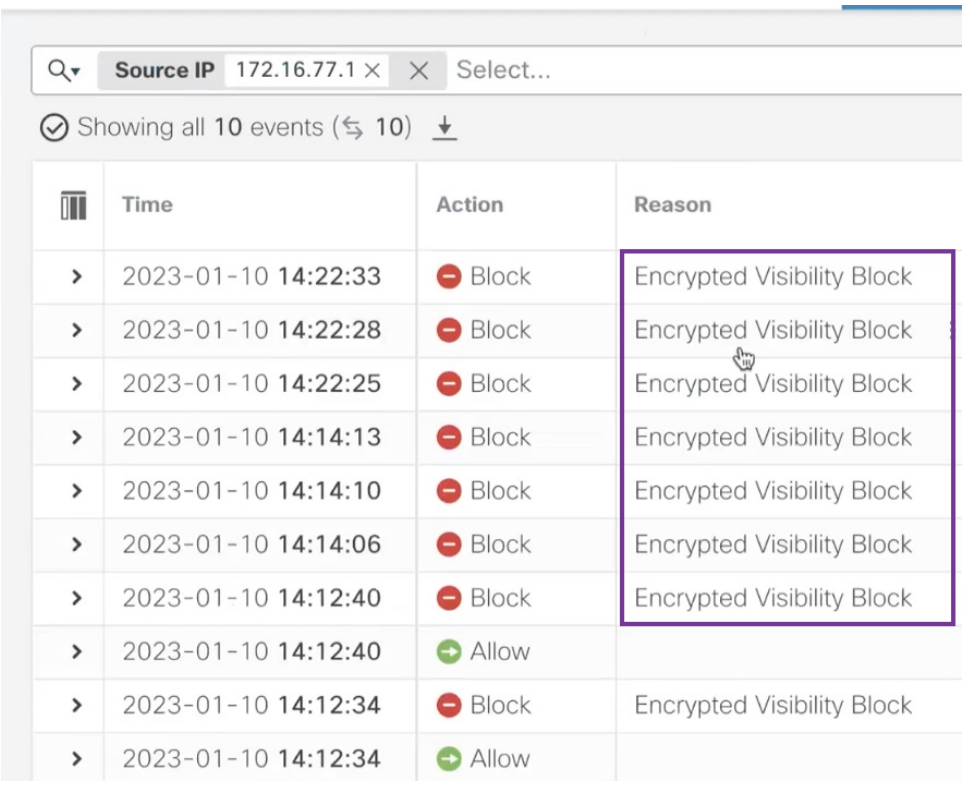
Deploy configuration changes.

View EVE Events

Procedure

Step 1 To verify the block action, choose **Analysis** > **Connections** > **Events**. You can also view the events from the **Unified Events** viewer.

Step 2 If you have configured EVE to block traffic, the **Reason** field shows **Encrypted Visibility Block**.



Step 3 The following is an example of the **Encrypted Visibility Process Name** as **test_malware**, **Encrypted Visibility Threat Confidence** as **Very High**, and **Encrypted Visibility Threat Confidence Score** as **90** percent.



Additional References

For detailed conceptual information, see the Encrypted Visibility Engine for Snort 3 chapter in this guide or the content in the following link:

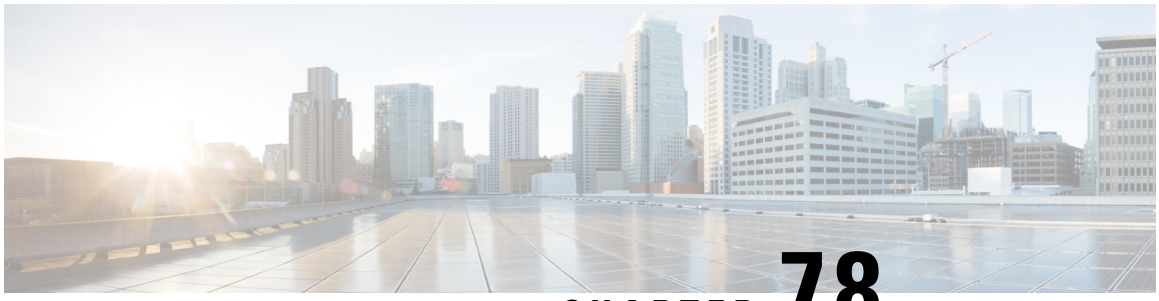
Encrypted Visibility Engine



PART XIX

Identity Policies for Access Control

- Identity Overview, on page 2255
- Realms, on page 2273
- Identity Source: Captive Portal, on page 2351
- Identity Source: ISE/ISE-PIC, on page 2373
- Identity Source: Passive Identity Agent, on page 2403
- Identity Source: pxGrid Cloud Identity (ISE 3.3 and Earlier), on page 2433
- Identity Source: pxGrid Cloud Identity (ISE 3.4 and Later), on page 2465
- Identity Source: Remote Access VPN, on page 2495
- Identity source: TS agent, on page 2501
- Identity Policies, on page 2505



CHAPTER 78

Identity Overview

The following topics discuss user identity:

- User identity, on page 2255
- Recommendation: identity realm limit, on page 2269
- Cloud-Delivered Firewall Management Center host and user limits, on page 2269
- User limits for microsoft azure active directory realms, on page 2271

User identity

User identity information can help you to identify the source of policy breaches, attacks, or network vulnerabilities, and trace them to specific users.

User identity information can help you determine:

- Who owns the host targeted by an intrusion event that has a Vulnerable (level 1: red) impact level.
- Who initiated an internal attack or portscan.
- Who is attempting unauthorized access to a specified host.
- Who is consuming an unreasonable amount of bandwidth.
- Who has not applied critical operating system updates.
- Who is using instant messaging software or peer-to-peer file-sharing applications in violation of company policy.
- Who is associated with each indication of compromise on your network.

Armed with this information, you can use other features of the system to mitigate risk, perform access control, and take action to protect others from disruption. These capabilities also significantly improve audit controls and enhance regulatory compliance.

After you configure user identity sources to gather user data, you can perform user awareness and user control.

For more information about identity sources, see About user identity sources, on page 2256.

Identity terminology

Identity terminology is a set of concepts that

- identifies users on networks using identity sources
- enables user awareness from authoritative and non-authoritative sources, and
- manages user access using identity policies.

Key identity terminology

This topic discusses common terminology for user identity and user control.

- **User awareness**: Identifying users on your network using *identity sources* (such as, TS Agent, or Cisco ISE). User awareness enables you to identify users from both *authoritative* (such as Active Directory) and *non-authoritative* (application-based) sources. To use Active Directory as an identity source, you must configure a realm and directory. For more information, see About user identity sources, on page 2256.
- **User control**: Configuring an *identity policy* that you associate with an *access control policy*. (The identity policy is then referred to as an access control *subpolicy*.) The identity policy specifies the identity source and, optionally, users and groups belonging to that source. By associating the identity policy with an access control policy, you determine whether to monitor, trust, block, or allow users or user activity in traffic on your network. For more information, see Access Control Policies, on page 1997.
- **Authoritative identity source**: a trusted server that validates user logins (for example, Active Directory). You can use the data obtained from authoritative logins to perform user awareness and user control. Authoritative user logins are obtained from passive and active authentications:
  - *Passive authentications* occur when a user authenticates through an external repository. ISE/ISE-PIC, the TS Agent, Microsoft Active Directory, and Microsoft Azure Active Directory, and the Passive Identity Agent are passive authentication user repositories supported by the system.

    Passive authentications occur when a user authenticates with Cisco ISE.
  - *Active authentications* occur when a user authenticates through preconfigured managed devices. Captive portal is another name for active authentication. Active authentication generally uses the same user repositories as passive authentication (the exceptions being ISE/ISE-PIC, and TS Agent, and the Passive Identity Agent, which are passive only).
- **Non-authoritative identity source**: an unknown or untrusted server that validates the user login. Traffic-based detection is the only non-authoritative identity source supported by the system. You can use the data obtained from non-authoritative logins to perform user awareness.

About user identity sources

Learn how to select and configure user identity sources for the firewall to enable user awareness and access control policies. This overview covers supported identity sources, authentication types, and key considerations for optimizing user identity management.

The table provides a brief overview of the user identity sources supported by the system. Each identity source provides a store of users for user awareness. These users can then be controlled with identity and access control policies.

User Identity Source	Server Requirements	Login Type	Authentication Type	User Control	For more, see...
Captive portal	OpenLDAP Microsoft Active Directory Microsoft Azure Active Directory	Authoritative	Active	Yes	Captive portal identity source, on page 2351 Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal), on page 2310
Passive authentication	OpenLDAP Microsoft Active Directory	Non-authoritative	Active	Yes	Create an LDAP realm or an Active Directory realm and realm directory, on page 2325
Passive authentication	Microsoft Azure Active Directory	Passive	Passive	Yes	Configure Microsoft Entra ID for passive authentication, on page 2293
Passive authentication with the TS Agent	Microsoft Windows Terminal Server	Authoritative	Passive	Yes	Terminal services agent identity source, on page 2501
Passive authentication with pxGrid Cloud identity source	Cisco ISE	Non-authoritative	Passive	Yes	pxGrid Cloud identity source, on page 2433
Remote Access VPN	OpenLDAP or Microsoft Active Directory	Authoritative	Active	Yes	Remote access VPN identity sources, on page 2495
	RADIUS	Authoritative	Active	No, awareness only	
ISE/ISE-PIC	Microsoft Active Directory	Authoritative	Passive	Yes	The ISE/ISE-PIC identity source, on page 2373
Traffic-based detection (Configured in the network discovery policy.)	—	Non-authoritative	—	No, awareness only	Traffic-based detection identity sources

Consider the information when selecting identity sources to deploy:

- You must use traffic-based detection for non-LDAP user logins.

- You must use traffic-based detection or captive portal to record failed login or authentication activity. A failed login or authentication attempt does not add a new user to the list of users in the database.
- The captive portal identity source requires a managed device with a routed interface. You *cannot* use an inline (also referred to as tap mode) interface with captive portal.

Data from those identity sources is stored in the Security Cloud Control's users database and the user activity database. You can configure Security Cloud Control-server user downloads to automatically and regularly download new user data to your databases.

After you configure identity rules using the desired identity source, you must associate each rule with an access control policy and deploy the policy to managed devices for the policy to have any effect. For more information about access control policies and deployment, see Associate prefilter, decryption, and identity policies with an access control, on page 2016.

For general information about user identity, see User identity, on page 2255.

Best practices for user identity

General recommendations for identity policy setup

Review the following recommendations before you set up identity policies.

- Know user limits
- Health monitor
- Use latest version of ISE/ISE-PIC, two types of remediation
- Captive portal requires routed interface, several individual tasks

Microsoft active directory and LDAP

Create one realm per LDAP server or Active Directory domain. The system supports Active Directory, LDAP, and other user repositories for user awareness and control. The association between an Active Directory or LDAP repository and the Security Cloud Control is referred to as a *realm*. For details about which versions are supported, see Supported servers for realms, on page 2278.

- The only user identity source supported by LDAP is captive portal. To use other identity sources (with the exception of ISE/ISE-PIC), you must use Active Directory.
- For Active Directory only: Create one *directory* per domain controller. For details, see Create an LDAP realm or an Active Directory realm and realm directory, on page 2325
- Users and groups in trust relationships between two domains are supported provided you add all Active Directory domains and domain controllers as realms and directories, respectively. For more information, see Realms and trusted domains, on page 2275.

Passive Identity Agent identity source

The passive identity agent software can monitor multiple Microsoft Active Directory servers and domain controllers and send user name and IP address information to the Security Cloud Control. For more information, see The passive identity agent identity source, on page 2403.

pxGrid Cloud Identity Source

The Cisco Identity Services Engine (Cisco ISE) pxGrid Cloud identity source enables you to use subscription and user data from Cisco ISE in Cloud-Delivered Firewall Management Center access control rules. Also, the identity source uses constantly changing dynamic objects from Cisco ISE in access control policies in the Cloud-Delivered Firewall Management Center.

For more information, see pxGrid Cloud identity source, on page 2433.

Proxy sequence

A *proxy sequence* is one or more managed devices that can be used to communicate with an LDAP, Active Directory, or ISE/ISE-PIC server. It is necessary only if Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.)

Although you can use one managed device as a proxy sequence, we strongly recommend you set up two or more so that, in the event one managed device cannot communicate with Active Directory or ISE/ISE-PIC, another managed device can take over.

Health monitor

The Security Cloud Control health monitor provides valuable information about the status of various Security Cloud Control functions, including user/realm mismatches, Snort memory usage, and ISE connection status.

- For more information about health modules, see *Health Modules* in the *Cisco Secure Firewall Management Center Administration Guide*.
- To set up policies to monitor health modules, see *Creating Health Policies* in the *Cisco Secure Firewall Management Center Administration Guide*.

Device-specific user limits

- User limits are discussed in User limits for microsoft active directory.
- If you use the ISE/ISE-PIC identity source, you can optionally limit the subnets the Cloud-Delivered Firewall Management Center monitors to reduce memory usage using identity mapping filters as discussed in Create an identity policy, on page 2507.

Use the latest version of ISE/ISE-PIC

If you expect to use the ISE/ISE-PIC identity source, we strongly recommend you always use the latest version to make sure you get the latest features and bug fixes.

- pxGrid 2.0 (which is used by version 2.6 patch 6 or later; or 2.7 patch 2 or later) also changes the remediation used by ISE/ISE-PIC from Endpoint Protection Service (EPS) to Adaptive Network Control (ANC). If you upgrade ISE/ISE-PIC, you must migrate your mediation policies from EPS to ANC.
- More information about using ISE/ISE-PIC can be found in Best practices for ISE/ISE-PIC integration, on page 2376.
- To set up the ISE/ISE-PIC identity source, see ISE/ISE-PIC configurations for user control , on page 2379.

Captive portal information

You can use captive portal active authentication with any of the following: LDAP, Microsoft AD, or Microsoft Azure AD (SAML).

- For LDAP, see Captive portal identity source, on page 2351.
- For Microsoft AD, see Captive portal identity source, on page 2351.
- For Microsoft Azure AD (SAML), see Create a Microsoft Azure AD (SAML) realm, on page 2289.

TS agent information

The TS Agent user identity source is required to identify user sessions on a Windows Terminal Server. The TS Agent software must be installed on the Terminal Server machine as discussed in the *Cisco Terminal Services (TS) Agent Guide*. In addition, you must synchronize the time on your TS Agent server with the time on the Security Cloud Control.

- TS Agent data is visible in the Users, User Activity, and Connection Event tables and can be used for user awareness and user control.
- For more information, see Best practices for TS agent configuration, on page 2502.

Associate the identity policy with an access control policy

After you configure your realm, directory, and user identity source, you must set up identity rules in an identity policy. To make the policy effective, you must associate the identity policy with an access control policy.

- For more information about creating an identity policy, see Create an identity policy, on page 2507.
- For more information about creating identity rules, see Create an identity rule, on page 2515.
- To associate an identity policy with an access control policy, see Associate prefilter, decryption, and identity policies with an access control, on page 2016.

Identity deployments

The Secure Firewall Management Center uses various identity sources, such as TS Agent, ISE/ISE-PIC, pxGrid Cloud, passive identity agent, and captive portal, to collect and manage user and group data for access control and user awareness.

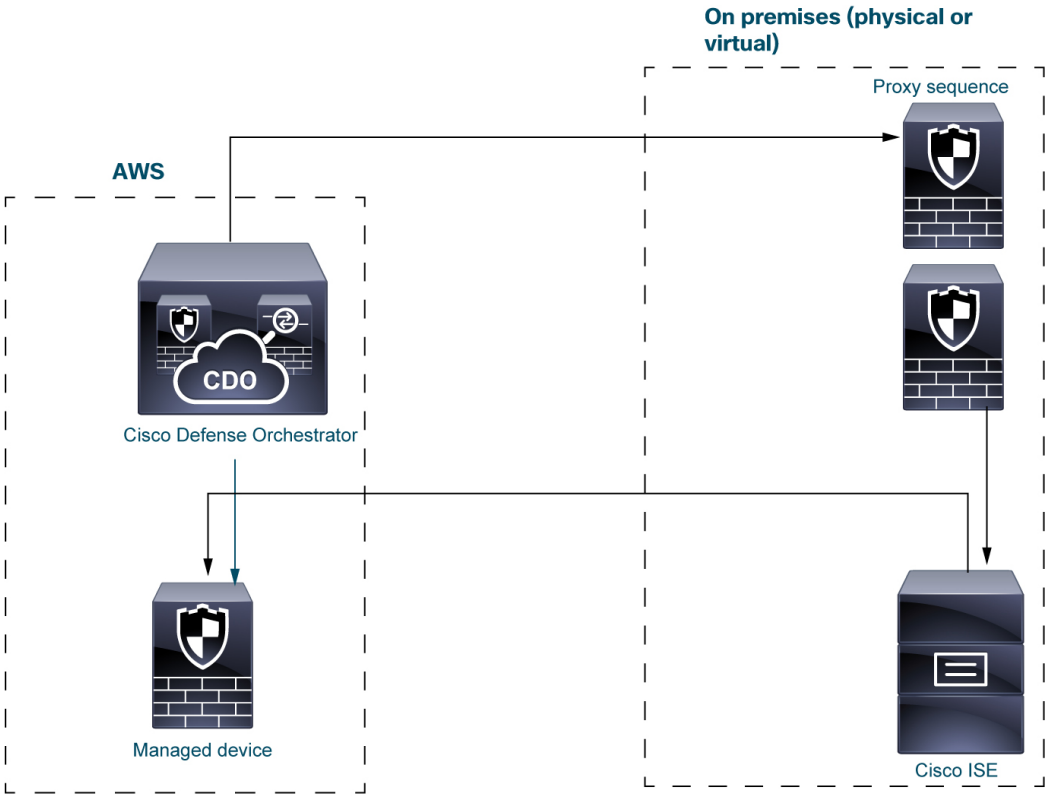
User data collection and storage

When the system detects user data from a user login, from any identity source, the user from the login is checked against the list of users in the Security Cloud Control user database. If the login user matches an existing user, the data from the login is assigned to the user. Logins that do not match existing users cause a new user to be created, unless the login is in SMTP traffic. Non-matching logins in SMTP traffic are discarded.

The group to which the user belongs is associated with the user as soon as the user is seen by the Security Cloud Control.

Here are examples of identity deployments.

The sample deployments discussed in this section are based on the system shown in this figure.



In the preceding figure, Security Cloud Control and one managed device are deployed to AWS and the other devices are located on premises. These devices can be either physical or virtual; they just need to be able to communicate with each other.

The two on-premises managed devices are intended to be used as a proxy sequence. You must add those devices to Security Cloud Control as well.

A *proxy sequence* is one or more managed devices that can be used to communicate with an LDAP, Active Directory, or ISE/ISE-PIC server. It is necessary only if Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.)

LDAP or Active Directory are required only for TS Agent and captive portal, as the following paragraphs explain.

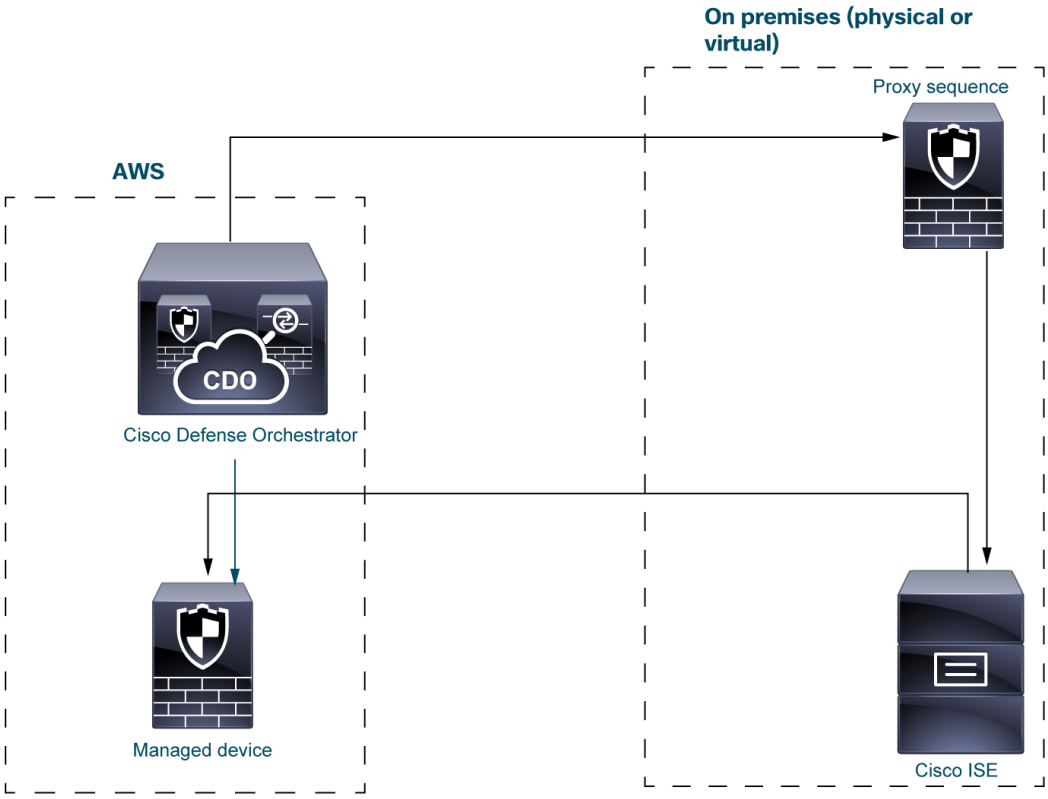
For more information about setting up a system like this, see Set up an identity policy, on page 2263.

ISE/ISE-PIC identity source

When you deploy the ISE/ISE-PIC identity source, Security Cloud Control contacts the proxy sequence if Security Cloud Control cannot contact the ISE/ISE-PIC server directly. Users, groups, and subscriptions are sent from the ISE/ISE-PIC server to the managed device in AWS.

You can optionally have an LDAP server in an ISE/ISE-PIC deployment but because it's optional, it isn't shown in this figure.

For more information about ISE/ISE-PIC, see The ISE/ISE-PIC identity source, on page 2373.



Passive identity agent source

Software that monitors one or several Microsoft AD domain controllers, sending user name and IP address information to the Security Cloud Control.

For more information, see The passive identity agent identity source, on page 2403.

pxGrid cloud identity source

pxGrid Cloud Identity Source

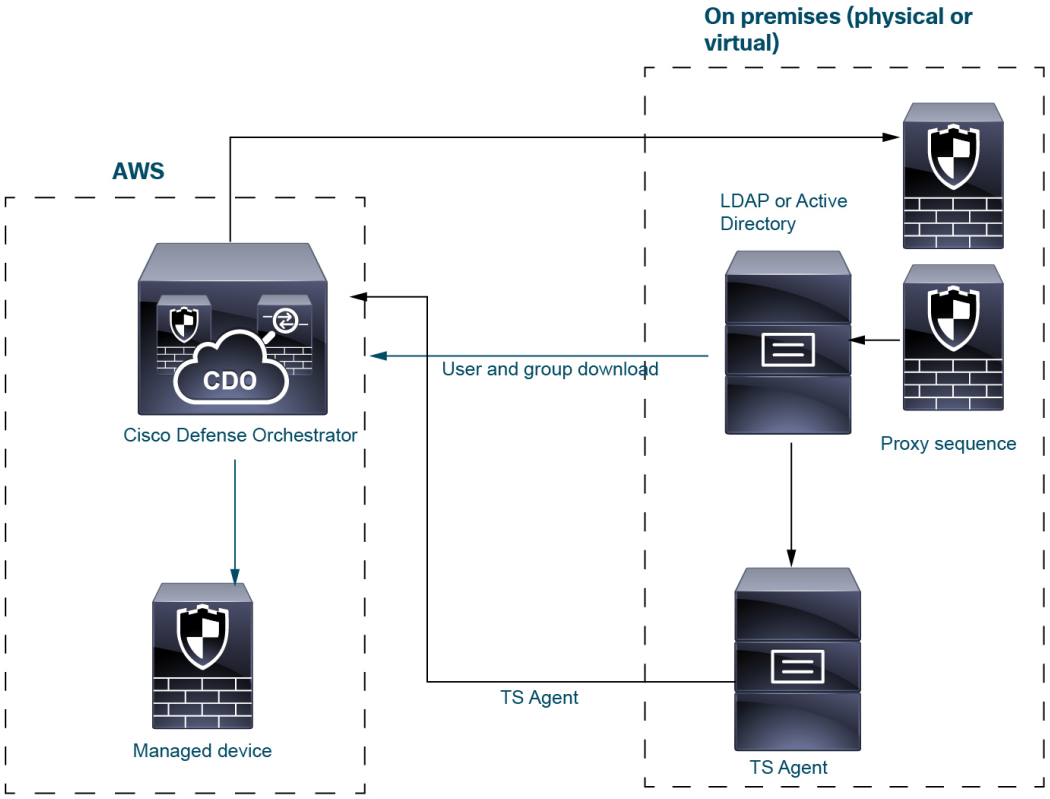
The Cisco Identity Services Engine (Cisco ISE) pxGrid Cloud identity source enables you to use subscription and user data from Cisco ISE in Cloud-Delivered Firewall Management Center access control rules. Also, the identity source uses constantly changing dynamic objects from Cisco ISE in access control policies in the Cloud-Delivered Firewall Management Center.

For more information, see pxGrid Cloud identity source, on page 2433.

TS Agent identity source

The Terminal Services (TS) Agent software runs on a Microsoft Server and sends Security Cloud Control user information based on the port range the users log in to the server with. TS Agent gets user identity information from LDAP or Active Directory and sends it to Security Cloud Control.

For more information about the TS Agent identity source, see Terminal services agent identity source, on page 2501.



Captive portal

You can use captive portal active authentication with any of these

- LDAP

  For more information, see Captive portal identity source, on page 2351.

- Microsoft AD

  For more information, see Captive portal identity source, on page 2351.

- Microsoft Azure AD (SAML)

  For more information, see Create a Microsoft Azure AD (SAML) realm, on page 2289.

Set up an identity policy

Set up an identity policy to control user access and filter network traffic using various user identity sources such as TS Agent, ISE/ISE-PIC, Microsoft Azure AD (SAML), captive portal.

An identity policy is not required if all these conditions are true:

- You use the ISE or ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies.

This topic provides a high-level overview of setting up an identity policy using any available user identity source: TS Agent, ISE/ISE-PIC, Microsoft Azure AD (SAML), captive portal.

Procedure

	Command or Action	Purpose
Step 1	Choose your identity source.	• To use Microsoft Active Directory or LDAP, continue with the next step. • To use a Microsoft Azure AD (SAML) realm, see Create a Microsoft Azure AD (SAML) realm, on page 2289. • To use the Cisco ISE identity source, see Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source, on page 2390. • To use the passive identity agent identity source, see The passive identity agent identity source, on page 2403. • To use the pxGrid Cloud identity source, see pxGrid Cloud identity source, on page 2433.
Step 2	(Optional.) Create a proxy sequence.	A *proxy sequence* is one or more managed devices that can be used to communicate with an LDAP, Active Directory, or ISE/ISE-PIC server. It is necessary only if Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.) Although you can use one managed device as a proxy sequence, we strongly recommend you set up two or more so that, in the event one managed device cannot communicate with Active Directory or ISE/ISE-PIC, another managed device can take over. See Create a proxy sequence, on page 2288.
Step 3	(Optional.) Create a realm and directory, one realm for every domain in the forest that contain users you want to use in user control. Also create one directory for every domain controller. Only users and groups that have corresponding Cloud-Delivered Firewall Management Center realms and directories can be used in identity policies.	Creating a realm, realm directory, and proxy sequence is optional if any of these are true: • You use SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.

	Command or Action	Purpose
		• You are using an identity policy only to filter network traffic. • A proxy sequence is required only if you use Security Cloud Control (Security Cloud Control) and it cannot directly communicate with Active Directory or ISE/ISE-PIC. The *realm* is a trusted user and group store, typically a Microsoft Active Directory repository. The Cloud-Delivered Firewall Management Center downloads users and groups at intervals you specify. You can include or exclude users and groups from being downloaded. See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325. For details about the options to create a realm, see Realm fields, on page 2281. A *directory* is an Active Directory domain controller that organizes information about a computer network's users and network shares. An Active Directory controller provides Directory Services for the realm. Active Directory distributes user and group objects across multiple domain controllers, which are peers that propagate local changes between each other by the use of Directory Services. For more information, see the Active Directory technical specification glossary on MSDN. You can specify more than one directory for a realm, in which case each domain controller is queried in the order listed on the realm's **Directory** tab page to match user and group credentials for user control. **Note** Configuring a realm or realm sequence is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.
Step 4	Synchronize users and groups from the realm.	To be able to control users and groups, you must synchronize them with the Cloud-Delivered Firewall Management Center. You can synchronize them with users and groups whenever you want or you can configure the

	Command or Action	Purpose
		system to synchronize them at a specified interval. When you synchronize users and groups, you can specify exceptions; for example, you can exclude the Engineering group from all user control for that realm, or you can exclude the user **`joe.smith`** from user controls that apply to the Engineering group. See Synchronize users and groups, on page 2330
Step 5	(Optional.) Create a realm sequence.	A realm sequence is an ordered list of realms that, when used in an identity policy, causes the system to search the realms in the specified order to find users to match the rule. See Create a realm sequence, on page 2332.
Step 6	Create a method to retrieve user and group data (the *identity source*).	Set up an identity source with its unique configuration to be able to control users and groups using data stored in the realm. Identity sources include TS Agent, captive portal, or Remote VPN. See one of these: • Configure the captive portal for user control, on page 2355 • Configure Remote Access VPN for AD-authenticated users, on page 2496
Step 7	Create an identity policy.	An identity policy contains one or more identity rules, optionally organized in categories. See Create an identity policy, on page 2507. **Note** Configuring a realm or realm sequence is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions; or if you use your identity policy only to filter network traffic.
Step 8	Create one or more identity rules.	Identity rules enable you to specify a number of matching criteria, including the type of authentication, network zones, networks or geolocation, realms, realm sequences, and so on. See Create an identity rule, on page 2515.
Step 9	Associate your identity policy with an access control policy.	An access control policy filters and optionally inspects traffic. An identity policy must be associated with an access control policy to have any effect. See Associate prefilter, decryption,

	Command or Action	Purpose
		and identity policies with an access control, on page 2016.

What to do next

1. Deploy the access control policy to at least one managed device.

   To use your policy to control user activity, the policy must be deployed to the managed devices to which clients connect. See Deploy configuration changes, on page 202.

2. Monitor user activity.

   View a list of active sessions collected by user identity sources or a list of user information collected by user identity sources. .

User activity database

The user activity database on the Cloud-Delivered Firewall Management Center is a logging system that contains records of user activity on your network detected or reported by all of your configured identity sources.

User activity logging events

The system logs events in these circumstances:

- When it detects individual logins or logoffs.
- When it detects a new user.
- When a system administrator manually delete a user.
- When the system detects a user that is not in the database, but cannot add the user because you have reached your user limit.
- When you resolve an indication of compromise associated with a user, or enable or disable indication of compromise rules for a user.

Note If the TS Agent monitors the same users as another passive authentication identity source (such as the ISE/ISE-PIC), the Cloud-Delivered Firewall Management Center prioritizes the TS Agent data. If the TS Agent and another passive source report identical activity from the same IP address, only the TS Agent data is logged to the Cloud-Delivered Firewall Management Center.

You can view user activity detected by the system using the Cloud-Delivered Firewall Management Center. (**Analysis** > **Users heading** > **User Activity**.)

Users database

A users database on the Cloud-Delivered Firewall Management Center is a system component that

- contains a record for each user detected or reported by all configured identity sources
- stores data obtained from authoritative sources for user control, and
- prioritizes user data based on identity source type when user limits are reached.

User database management and limitations

For more information about the supported non-authoritative and authoritative identity sources, refer to About user identity sources, on page 2256.

The total number of users the Cloud-Delivered Firewall Management Center can store depends on the Cloud-Delivered Firewall Management Center model. After the user limit is reached, the system prioritizes previously-undetected user data based on its identity source, as described in this list:

- Non-authoritative identity source: If the new user is from a non-authoritative identity source, the system does not add the user to the database. To allow new users to be added, you must delete users manually or with a database purge.
- Authoritative identity source: If the new user is from an authoritative identity source, the system deletes the non-authoritative user who has remained inactive for the longest period and adds the new user to the database.

If an identity source is configured to exclude specific user names, user activity data for those user names are not reported to the Cloud-Delivered Firewall Management Center. These excluded user names remain in the database, but are not associated with IP addresses.

If you have Cloud-Delivered Firewall Management Center high availability configured and the primary fails, no logins reported by a captive portal, ISE/ISE-PIC, TS Agent, or Remote Access VPN device can be identified during failover downtime, even if the users were previously seen and downloaded to the Cloud-Delivered Firewall Management Center. The unidentified users are logged as Unknown users on the Cloud-Delivered Firewall Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.



Note If the TS Agent monitors the same users as another passive authentication identity source (ISE/ISE-PIC), the Cloud-Delivered Firewall Management Center prioritizes the TS Agent data. If the TS Agent and another passive source report identical activity from the same IP address, only the TS Agent data is logged to the Cloud-Delivered Firewall Management Center.

When the system detects a new user session, the user session data remains in the users database until one of these conditions occurs:

- A user on the Cloud-Delivered Firewall Management Center manually deletes the user session.
- An identity source reports the logoff of that user session.
- A realm ends the user session as specified by the realm's **User Session Timeout: Authenticated Users**, **User Session Timeout: Failed Authentication Users**, or **User Session Timeout: Guest Users** setting.

Recommendation: identity realm limit

We limit you to a maximum of 25 realms regardless of Firewall Threat Defense model; the same limit applies to the Cloud-Delivered Firewall Management Center.

The number of *users* you can download to a system is limited according to the information provided in:

Cloud-Delivered Firewall Management Center user limit, on page 2270

Cloud-Delivered Firewall Management Center host and user limits

Cloud-Delivered Firewall Management Center host and user limits are restrictions that define the maximum number of hosts and users that can be managed or supported by the system.

User and host limit details

These topics describe the user and host limits for the Cloud-Delivered Firewall Management Center.

Cloud-Delivered Firewall Management Center host limit

Cloud-Delivered Firewall Management Center host limit describes the storage capacity that

- allows the system to store a maximum of 600,000 hosts in its host database
- automatically adds a host to the network map when activity is detected for an IP address in your monitored network, and
- controls host visibility and contextual data access based on network map inclusion.

Host limit recommendations

The system can store up to 600,000 hosts but we recommend limits based on the number of devices managed.

Number of devices managed by Security Cloud Control	Recommended number of hosts
1-50	100,000
51-300	300,000
301-1000	600,000

You cannot view contextual data for hosts not in the network map. However, you can perform access control. For example, you can perform application control on traffic to and from a host not in the network map, even though you cannot use a compliance allow list to monitor the host's network compliance.

Note The system counts MAC-only hosts separately from hosts identified by both IP addresses and MAC addresses. All IP addresses associated with a host are counted together as one host.

Reaching the Host Limit and Deleting Hosts

When you reach the host limit, the network discovery policy controls what happens when you detect a new host. You can drop the new host, or replace the host that has been inactive for the longest time. You can also set the period after which the system removes a host from the network map due to inactivity. Although you can manually delete a host, an entire subnet, or all of your hosts from the network map, if the system detects activity associated with a deleted host, it re-adds the host.

Cloud-Delivered Firewall Management Center user limit

A user is added to the Cloud-Delivered Firewall Management Center user database when:

- The user is downloaded from a realm.
- A captive portal or RA-VPN user logs in.
- A user is detected from any identity source (for example, TS Agent).

A Cloud-Delivered Firewall Management Center can store a maximum of 600,0000 users in its host database but we recommend these limits.

Number of devices managed by Cisco Security Cloud	Recommended number of users
1-50	100,000
51-300	300,000
301-1000	600,000

Note The Secure Firewall 200 series can download a *cumulative total* of 10,000 user IPs, SXP/SGT mappings, endpoint profiles, and dynamic objects. After a total of 10,000 have been downloaded, the Secure Firewall 200 stops downloading any objects until some previously downloaded objects have been removed. (For example, if users log out, the memory is free for other objects.) This limit applies regardless of how many devices Cisco Security Cloud manages.

Because of its limited capacity, we recommend you *do not* use a Secure Firewall 220 as a proxy sequence.

Only authoritative users are available for user control with access control policies.

The Cloud-Delivered Firewall Management Center can store 600,000 sessions in its user database.

When the system detects a new, previously-undetected user after the limit has been reached, it prioritizes user data based on their identity source:

- If the new user is from a non-authoritative source, the system does not add the non-authoritative user to the database. To allow new users to be added, you must delete users manually or purge the database.

- If the new user is from an authoritative identity source, the system deletes the non-authoritative user who has remained inactive for the longest period of and adds the new authoritative user to the database.

  If there are only authoritative users, the system deletes the authoritative user who has remained inactive for the longest period of time and adds the new user to the database.

Troubleshooting information can be found in Troubleshoot user control, on page 2526.



Tip Note that if you are using traffic-based detection, you can restrict user logging by protocol to help minimize username clutter and preserve space in the database. For example, you could prevent the system from adding users discovered in AIM, POP3, and IMAP traffic because you know it is traffic from specific contractors or visitors you do not want to monitor.

User limits for microsoft azure active directory realms

This reference provides user limit information for Microsoft Azure Active Directory realms, showing how different device models handle user monitoring and login capacity.

Microsoft azure active directory user limits

About user limits

Your Cloud-Delivered Firewall Management Center model determines how many individual users you can monitor.

Note the following:

- The maximum number of *downloaded* users depends on your Cloud-Delivered Firewall Management Center model.
- The maximum number of *concurrent* user sessions (that is, logins) depends on your managed device model. A single user can have multiple sessions from different unique IP addresses.



Note The system downloads all user sessions to all Firewall Threat Defense devices. If you have devices with different user concurrent user session limits, the Firewall Threat Defense with the smallest limit reports health warnings when its memory reaches the configured limit. (For example, if your Cloud-Delivered Firewall Management Center manages a Firepower 2110 and a 4125, the 2110 reports health warnings when the number of concurrent user sessions approaches its maximum of 64,000.)

Refer to these tables.

Table 127: Maximum Downloaded Users by Cloud-Delivered Firewall Management Center Model[1]

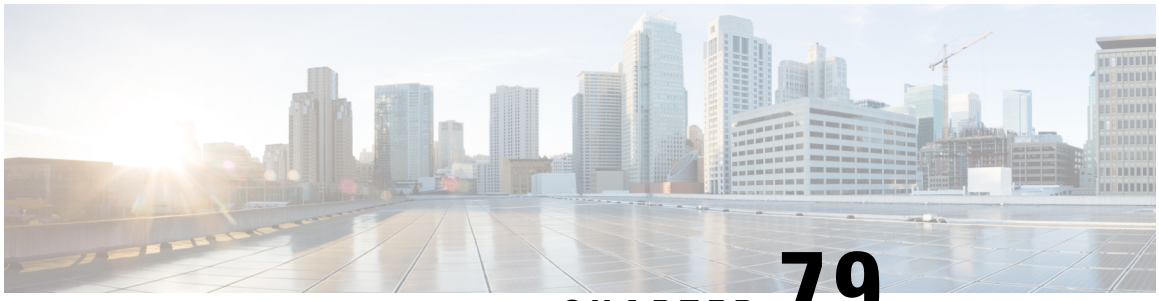
Cloud-Delivered Firewall Management Center Model	Number of Dynamic Attributes Connector Connectors	Maximum Downloaded Users
Firewall Management Center Virtual (any supported hypervisor)	10	50,000
Firewall Management Center Virtual 300 (any supported hypervisor)	20	150,000

[1]—Cloud-Delivered Firewall Management Center models are subject to end of life and end of sale. For more information, see End-Of-Life and End-Of-Sale Notices.

Table 128: Maximum Concurrent User Login Limits by Firewall Threat Defense

Firewall Threat Defense Model	Maximum Concurrent User Logins
Firewall Threat Defense Virtual 5, 10, 20, 30, 50 (any supported hypervisor)	50,000
Secure Firewall 220	10,000 (see note)
Firepower 1010, 1120, 1140, 1150 Firepower 2110, 2120, 2130 Secure Firewall 3110, 3120, 3130, 3140	50,000
Firepower 2140 Secure Firewall 3130, 3140	150,000
Firepower 4140, 4145, 4150 Firepower 9300 Secure Firewall 4215	225,000
Secure Firewall 4225, 4245	300,000

Note The Secure Firewall 200 series can download a *cumulative total* of 10,000 user IPs, SXP/SGT mappings, endpoint profiles, and dynamic objects. After a total of 10,000 have been downloaded, the Secure Firewall 200 stops downloading any objects until some previously downloaded objects have been removed. (For example, if users log out, the memory is free for other objects.)



CHAPTER 79

Realms

The following topics describe realms and identity policies:

- Realms and realm sequences, on page 2273
- License requirements for realms, on page 2280
- Requirements and prerequisites for realms, on page 2281
- Prerequisites for kerberos authentication, on page 2281
- Realm fields, on page 2281
- Realm directory and synchronize fields, on page 2285
- Create a proxy sequence, on page 2288
- Create a Microsoft Azure AD (SAML) realm, on page 2289
- Create a Microsoft Azure AD Realm, on page 2316
- Create an LDAP realm or an Active Directory realm and realm directory, on page 2325
- Create a realm sequence, on page 2332
- Cross-domain trust setups with Cloud-Delivered Firewall Management Center, on page 2333
- Manage a realm, on page 2341
- Compare realms, on page 2342
- Troubleshoot realms and user downloads, on page 2342
- History for realms, on page 2348

Realms and realm sequences

Realms are connections between the Cloud-Delivered Firewall Management Center and the user accounts on the servers you monitor that

- specify the connection settings and authentication filter settings for the server,
- specify the users and user groups whose activity you want to monitor, and
- query the user repository for user metadata on authoritative users, as well as some non-authoritative users: POP3 and IMAP users detected by traffic-based detection and users detected by traffic-based detection, a TS Agent, or ISE/ISE-PIC.

Realm configuration requirements

You can add multiple domain controllers as directories in a realm, but they must share the same basic realm information. The directories in a realm must be exclusively LDAP or exclusively Active Directory (AD)

servers. After you enable a realm, your saved changes take effect next time the Cloud-Delivered Firewall Management Center queries the server.

To perform user awareness, you must configure a realm for any of the supported server types. The system uses these connections to query the servers for data associated with POP3 and IMAP users, and to collect data about LDAP users discovered through traffic-based detection.

The system uses the email addresses in POP3 and IMAP logins to correlate with LDAP users on an Active Directory or OpenLDAP. For example, if a managed device detects a POP3 login for a user with the same email address as an LDAP user, the system associates the LDAP user's metadata with that user.

To perform user control, you can configure any of these:

- A realm or realm sequence for an Active Directory server, or for ISE/ISE-PIC



Note Configuring a Microsoft AD realm or realm sequence is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions; or if you use your identity policy only to filter network traffic.

- A realm or realm sequence for a Microsoft AD server for the TS Agent.
- For captive portal, an LDAP realm.

  A realm sequence is not supported for LDAP.

You can nest Microsoft AD groups and the Cloud-Delivered Firewall Management Center downloads those groups and the users they contain. You can optionally restrict which groups and users get downloaded as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

User synchronization

You can configure a realm or realm sequence to establish a connection between the Cloud-Delivered Firewall Management Center and an LDAP or Microsoft AD server to retrieve user and user group metadata for certain detected users:

- LDAP and Microsoft AD users authenticated by captive portal or reported by ISE/ISE-PIC. This metadata can be used for user awareness and user control.
- POP3 and IMAP user logins detected by traffic-based detection, if those users have the same email address as an LDAP or AD user. This metadata can be used for user awareness.

The Cloud-Delivered Firewall Management Center obtains this information and metadata about each user:

- LDAP user name
- First and last names
- Email address
- Department
- Telephone number

Important To reduce latency between the Cloud-Delivered Firewall Management Center and your Active Directory domain controller, we strongly recommend you configure a realm directory (that is, domain controller) that is as close as possible geographically to the Cloud-Delivered Firewall Management Center.

For example, if your Cloud-Delivered Firewall Management Center is in North America, configure a realm directory that is also in North America. Failure to do so can cause problems such as timeout downloading users and groups.

User activity data

User activity data is stored in the user activity database and user identity data is stored in the users database. If your access control parameters are too broad, the Cloud-Delivered Firewall Management Center obtains information on as many users as it can and reports the number of users it failed to retrieve in the Tasks tab page of the Message Center.

To optionally limit the subnets on which a managed device watches for user awareness data, you can use the **configure identity-subnet-filter** command as discussed in the *Cisco Secure Firewall Threat Defense Command Reference*.

Note If you remove a user that has been detected by the system from your user repository, the Cloud-Delivered Firewall Management Center does *not* remove that user from its users database; you must manually delete it. However, your LDAP changes *are* reflected in access control rules when the Cloud-Delivered Firewall Management Center next updates its list of authoritative users.

Realms and trusted domains

When you configure a Microsoft Active Directory (AD) *realm* in the Cloud-Delivered Firewall Management Center, it is associated with a Microsoft Active Directory or LDAP *domain*.

Microsoft active directory forests and trust relationships

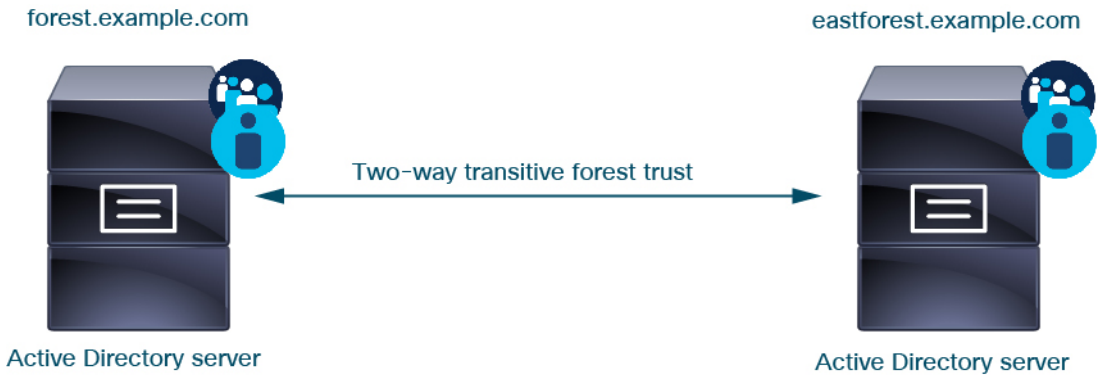
A grouping of Microsoft Active Directory (AD) domains that trust each other is commonly referred to as a *forest*. This trust relationship can enable domains to access each other's resources in different ways. For example, a user account defined in domain A can be marked as a member of a group defined in domain B.

The system supports AD forests that are configured in a trust relationship. There are several types of trust relationships; this guide discusses two-way, transitive forest trust relationships. This simple example shows two forests: **forest.example.com** and **eastforest.example.com**. Users and groups in each forest can be authenticated by AD in the other forest, provided you configure the forests that way.

If you set up the system with one realm for each domain and one directory for each domain controller, the system can discover up to 100,000 foreign security principals (users and groups). If these foreign security principals match a user downloaded in another realm, then they can be used in access control policy.

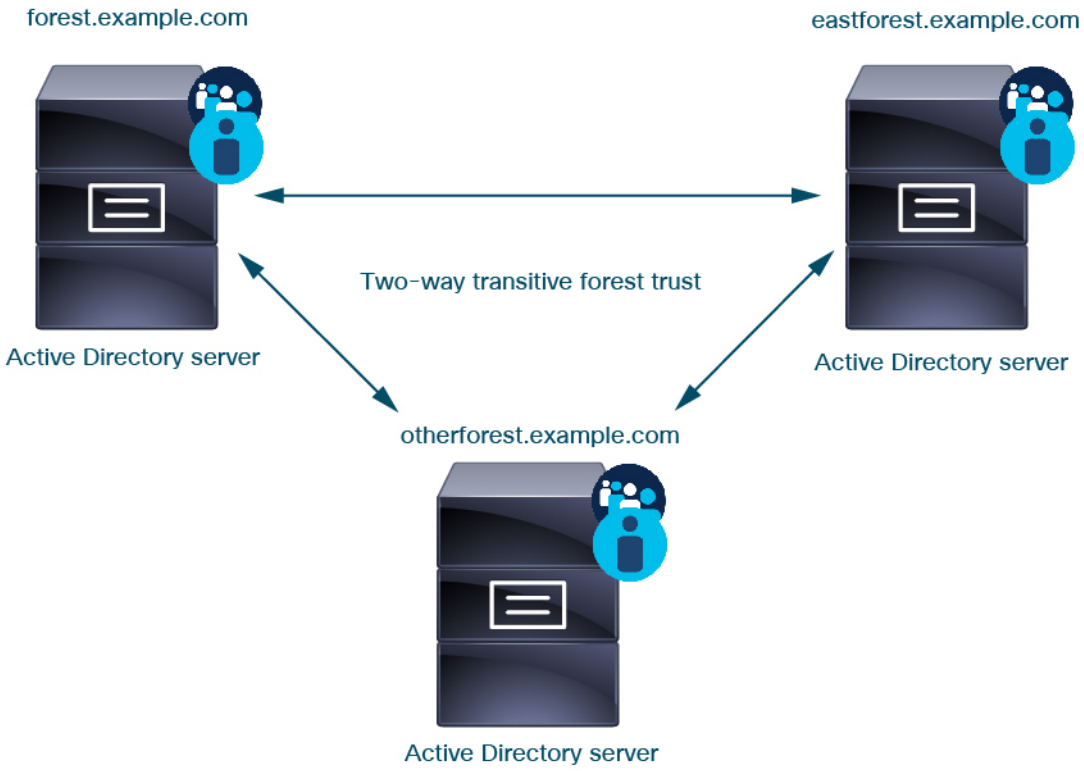
You need not configure a realm for any domain that has no users you wish to use in access control policies.

Figure 387: Two-forest configuration with two-way transitive trust



To continue the example, suppose you have three AD forests (one of which could be a subdomain or an independent forest), all set up as two-way transitive forest relationships, all users and groups are available in all three forests as well as in the system. (As in the preceding example, all three AD domains must be set up as realms and all domain controllers must be configured as directories in those realms.)

Figure 388: Three-forest configuration with two-way transitive trust

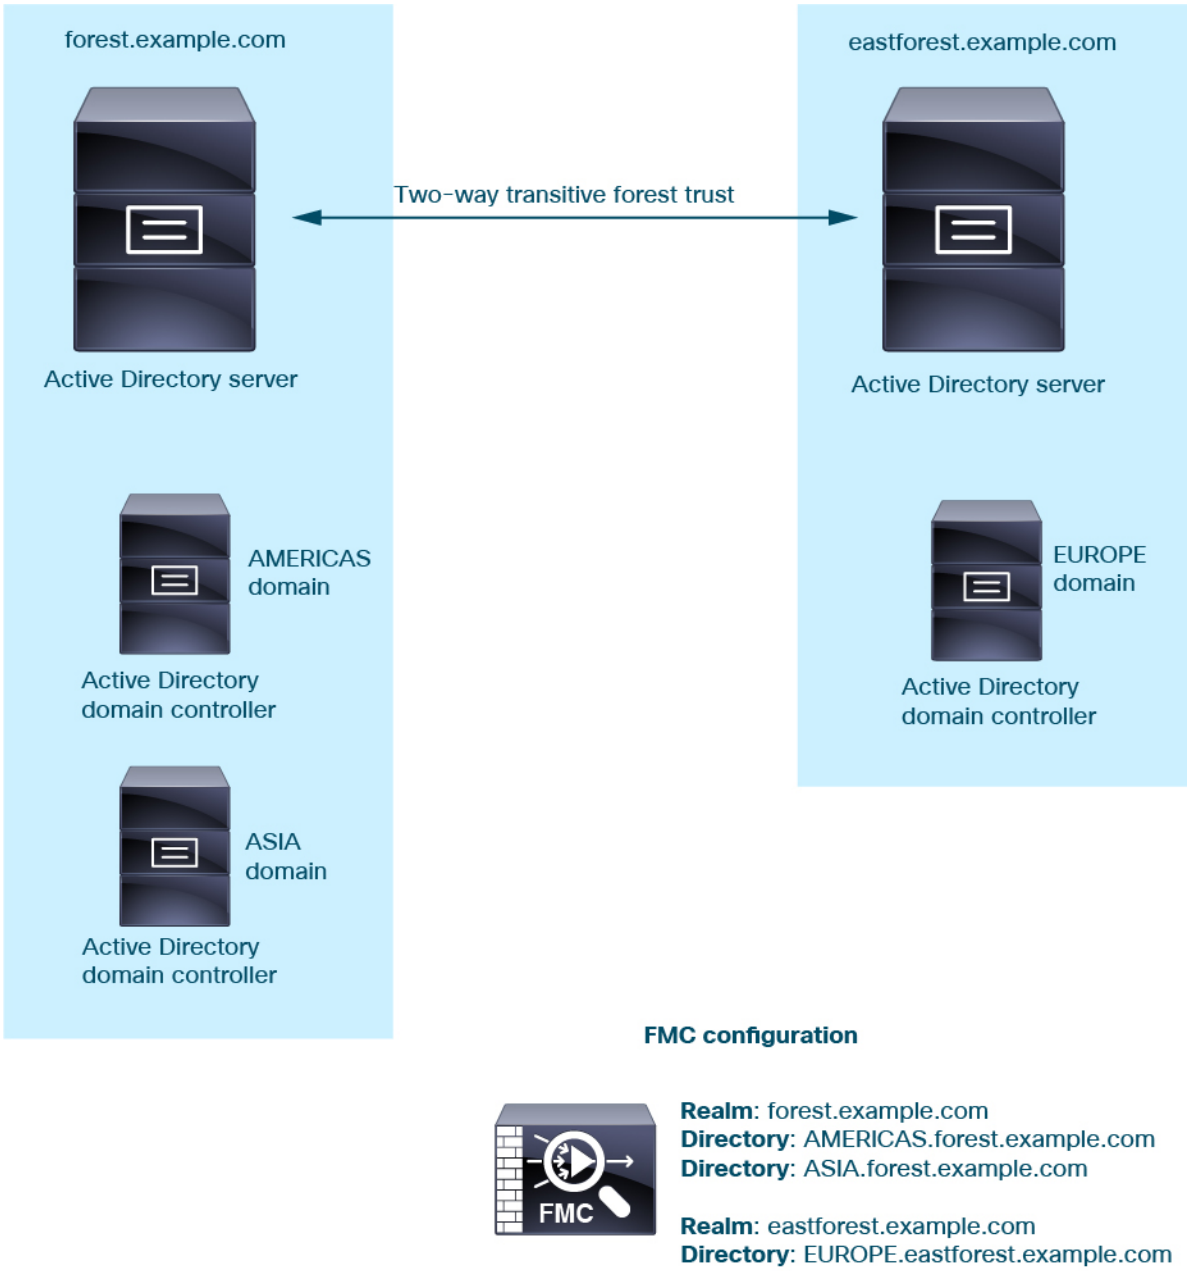


Finally, you can set up the Cloud-Delivered Firewall Management Center to be able to enforce identity policies on users and groups in a two-forest system with two-way transitive forest trust. Suppose each forest has at least one domain controller, each of which authenticates different users and groups. For the Cloud-Delivered Firewall Management Center to be able to enforce identity policies on those users and groups, you must set

up each domain containing relevant users as Cloud-Delivered Firewall Management Center realm and each domain controller as Cloud-Delivered Firewall Management Center directory in the respective realm.

Failure to properly configure the Cloud-Delivered Firewall Management Center prevents some of the users and groups from being able to be used in policies. You will see warnings when you try to synchronize users and groups in that case.

Figure 389: Multi-domain forest configuration



Using the preceding example, set up the Cloud-Delivered Firewall Management Center as follows:

- Realm for any domain in **forest.example.com** that contains users you want to control with access control policies
  - Directory in the realm for **AMERICAS.forest.example.com**
  - Directory in the realm for **ASIA.forest.example.com**
- Realm for any domain in **eastforest.example.com** that contains users you want to control with access control policies
  - Directory in the realm for **EUROPE.eastforest.example.com**

Note The Cloud-Delivered Firewall Management Center uses the AD field **msDS-PrincipalName** to resolve references to find user and group names in each domain controller. **msDS-PrincipalName** returns a NetBIOS name.

Supported servers for realms

You can configure realms to connect to the following types of servers, providing they have TCP/IP access from the Cloud-Delivered Firewall Management Center:

Server Type	Supported for ISE/ISE-PIC data retrieval?	Supported for TS Agent data retrieval?	Supported for captive portal data retrieval?
Microsoft Active Directory on Windows Server 2012, 2016, and 2019	Yes	Yes	Yes
OpenLDAP on Linux	No	No	Yes

An Active Directory Global Catalog server is *not supported* as a realm directory. For more information about the Global Catalog Server, refer to Global Catalog on learn.microsoft.com.

Note If the TS Agent is installed on a Microsoft Active Directory Windows Server shared with another passive authentication identity source (ISE/ISE-PIC), the Cloud-Delivered Firewall Management Center prioritizes the TS Agent data. If the TS Agent and a passive identity source report activity by the same IP address, only the TS Agent data is logged to the Cloud-Delivered Firewall Management Center.

Note the following about your server group configurations:

- To perform user control on user groups or on users in groups, you must configure user groups on the LDAP or Active Directory server.
- Group names cannot start with `s-` because it is used internally by LDAP.

Neither group names nor organizational unit names can contain special characters like asterisk (*), equals (=), or backslash (\); otherwise, users in those groups or organizational units are not downloaded and are not available for identity policies.

- To configure an Active Directory realm that includes or excludes users who are members of a sub-group on your server, note that Microsoft recommends that Active Directory has no more than 5000 users per group in Windows Server 2012. For more information, see Active Directory Maximum Limits—Scalability on MSDN.

  If necessary, you can modify your Active Directory server configuration to increase this default limit and accommodate more users.

- To uniquely identify the users reported by a server in your Remote Desktop Services environment, you must configure the Cisco Terminal Services (TS) Agent. When installed and configured, the TS Agent assigns unique ports to individual users so the system can uniquely identify those users. (Microsoft changed the name *Terminal Services* to *Remote Desktop Services*.)

  For more information about the TS Agent, see the *Cisco Terminal Services (TS) Agent Guide*.

Supported server object class and attribute names

Map user metadata attributes from your LDAP servers to the correct firewall fields to ensure user information is accurately retrieved and populated in the firewall.

The servers in your realms *must* use the attribute names listed in the following table for the Cloud-Delivered Firewall Management Center to retrieve user metadata from the servers. If the attribute names are incorrect on your server, the Cloud-Delivered Firewall Management Center cannot populate its database with the information in that attribute.

Table 129: Map of attribute names to Cloud-Delivered Firewall Management Center fields

Metadata	Cloud-Delivered Firewall Management Center Attribute	LDAP ObjectClass	Active Directory Attribute	OpenLDAP Attribute
LDAP user name	Username	• user • inetOrgPerson	samaccountname	cn uid
first name	First Name		givenname	givenname
last name	Last Name		sn	sn
email address	Email		mail userprincipalname (if mail has no value)	mail
department	Department		department distinguishedname (if department has no value)	ou
telephone number	Phone		telephonenumber	telephonenumber

Note The LDAP ObjectClass for groups is `group`, `groupOfNames`, (`group-of-names` for Active Directory) or `groupOfUniqueNames`.

For more information about ObjectClasses and attributes, see these references:

- Microsoft Active Directory:
  - ObjectClasses: All Classes on MSDN
  - Attributes: All Attributes on MSDN
- OpenLDAP: RFC 4512

License requirements for realms

This topic outlines the specific license requirements necessary to enable and manage realms.

Threat Defense License

Any

Classic License

Control

Requirements and prerequisites for realms

Model support

Any.

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Prerequisites for kerberos authentication

If you're using Kerberos to authentication captive portal users, keep the following in mind.

Hostname character limit

If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.

DNS response character limit

DNS must return a response of 64KB or less to the hostname; otherwise, the AD connection test fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.

Realm fields

Configure a realm to manage user authentication, directory access, and session timeouts for Active Directory, LDAP, or LOCAL user repositories in the Secure Firewall Management Center. Learn how to set realm type, credentials, supported characters, and session duration for identity policies and RA-VPN configurations.

Realm configuration fields

These settings apply to all Active Directory servers or domain controllers (also referred to as *directories*) in a realm.

Name

A unique name for the realm.

- To use the realm in identity policies, the system supports alphanumeric and special characters.
- To use the realm in RA VPN configurations, the system supports alphanumeric, hyphen (-), underscore (_), and plus (+) characters.

Description

(Optional.) Enter a description of the realm.

Type

The type of realm, **AD** for Microsoft Active Directory, **LDAP** for other supported LDAP repositories, or **LOCAL**. For a list of supported LDAP repositories, see Supported servers for realms, on page 2278. You can authenticate captive portal users with an LDAP repository; all others require Active Directory.

Note Only captive portal supports an LDAP realm.

The realm type **LOCAL** is used for configuring LOCAL user settings. The LOCAL realm is used in remote access user authentication.

Add the following LOCAL User Information for the LOCAL realm:

- **Username**—Name of the LOCAL user.
- **Password**—LOCAL user password.
- **Confirm Password**—Confirm the LOCAL user password.

Note **Click Add another LOCAL user** to add more users to the LOCAL realm.

You can add more users after creating the realm and update password for the LOCAL users. You can also create multiple LOCAL realms but cannot disable them.

AD Primary Domain

For Microsoft Active Directory realms only. Domain for the Active Directory server where users should be authenticated.

Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different Microsoft AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

AD Join Username and AD Join Password

(Available on the **Realm Configuration** tab page when you edit a realm.)

For Microsoft Active Directory realms intended for Kerberos captive portal active authentication, the distinguished username and password of any Active Directory user with appropriate rights to create a Domain Computer account in the Active Directory domain.

Keep the following in mind:

- DNS must be able to resolve the domain name to an Active Directory domain controller's IP address.
- The user you specify must be able to join computers to the Active Directory domain.
- The user name must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).

If you choose **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) as the **Authentication Protocol** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.



Note The SHA-1 hash algorithm is not secure for storing passwords on your Active Directory server and should not be used. For more information, consult a reference such as Migrating your Certification Authority Hashing Algorithm from SHA1 to SHA2 on Microsoft TechNet or Password Storage Cheat Sheet on the Open Web Application Security Project website.

We recommend SHA-256 for communicating with Active Directory.

Group Attribute

(Available on the **Realm Configuration** tab page.)

These group attributes are available for Microsoft Active Directory or LDAP realms.

- Member: A user attribute that indicates the groups to which a user belongs. Each user entry contains a Member attribute that lists the groups the user is part of, rather than each group listing its members.
- Unique Member: A group attribute that lists the distinguished names (DNs) of the members of the group.

Directory Username and Directory Password

The distinguished username and password for a user with appropriate access to the user information you want to retrieve.

Note the following:

- For OpenLDAP, the user's access privileges are determined by the `<level>` parameter discussed in section 8 of the OpenLDAP specification. The user's `<level>` should be `auth` or better.
- The user name must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).

Note The SHA-1 hash algorithm is not secure for storing passwords on your Active Directory server and should not be used. For more information, consult a reference such as Migrating your Certification Authority Hashing Algorithm from SHA1 to SHA2 on Microsoft TechNet or Password Storage Cheat Sheet on the Open Web Application Security Project website.

We recommend SHA-256 for communicating with Active Directory.

Base DN

(Optional.)The directory tree on the server where the Security Cloud Control should begin searching for user data. If you don't specify a **Base DN**, the system retrieves the top-level DN provided you can connect to the server.

Typically, the base distinguished name (DN) has a basic structure indicating the company domain name and operational unit. For example, the Security organization of the Example company might have a base DN of `ou=security,dc=example,dc=com`.

Group DN

(Optional.)The directory tree on the server where the Security Cloud Control should search for users with the group attribute. A list of supported group attributes is shown in Supported server object class and attribute names, on page 2279. If you don't specify a **Group DN**, the system retrieves the top-level DN provided you can connect to the server.

Note The following is the list of characters the system *supports* in users, groups, DNS in your directory server. Using any characters other than these could result in the system failing to download users and groups.

Entity	Supported characters
User name	**a-z A-Z 0-9 ! # $ % ^ & ( ) _ - { } ' . ~ `**
Group name	**a-z A-Z 0-9 ! # $ % ^ & ( ) _ - { } ' . ~ `**
Base DN and Group DN	**a-z A-Z 0-9 ! @ $ % ^ & * ( ) _ - . ~ `**

A space is not supported anywhere in a user name, including at the end.

Proxy

From the list, click one or more managed devices or a proxy sequence. These devices must be able to communicate with Active Directory or ISE/ISE-PIC to retrieve user data for identity policies.

These fields are available when you edit an existing realm.

User Session Timeout

(Available on the **Realm Configuration** tab page when you edit a realm.)

Enter the number of minutes before user sessions time out. The default is 1440 (24 hours) after the user's login event. After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Security Cloud Control as Unknown (except for **Failed Captive Portal Users**).

In addition, if you set up ISE/ISE-PIC without a realm and the timeout is exceeded, a workaround is required. For more information, contact Cisco TAC.

You can set timeout values for the following:

- **User Agent and ISE/ISE-PIC Users**: Timeout for users tracked by the user agent or by ISE/ISE-PIC, which are types of passive authentication.

  The timeout value you specify does *not* apply to pxGrid SXP session topic subscriptions (for example, destination SGT mappings). Instead, session topic mappings are preserved as long as there is no delete or update message for a given mapping from ISE.

  For more information about ISE/ISE-PIC, see The ISE/ISE-PIC identity source, on page 2373.

- **Terminal Services Agent Users**: Timeout for users tracked by the TS Agent, which is a type of passive authentication. For more information, see Terminal services agent identity source, on page 2501.

- **Captive Portal Users**: Timeout for users who successfully log in using the captive portal, which is a type of active authentication. For more information, see Captive portal identity source, on page 2351.

- **Failed Captive Portal Users**: Timeout for users who do not successfully log in using the captive portal. You can configure the **Maximum login attempts** before the user is seen by the Security Cloud Control as Failed Auth User. A Failed Auth User can optionally be granted access to the network using access control policy and, if so, this timeout value applies to those users.

  For more information about failed captive portal logins, see Captive portal fields, on page 2367.

- **Guest Captive Portal Users**: Timeout for users who log in to the captive portal as a guest user. For more information, see Captive portal identity source, on page 2351.

Realm directory and synchronize fields

Configure Active Directory realm directory and user synchronization fields to enable secure authentication, user data lookup, and group management in the Secure Firewall Management Center. Learn how to specify server, encryption, domain, and group settings for optimized user awareness and access control.

Realm directory fields

These settings apply to individual servers (such as Active Directory domain controllers) in a realm.

Hostname / IP Address

Fully qualified host name of the Active Directory domain controller machine. To find the fully qualified name, see Find the active directory server's name, on page 2328.

If you're using Kerberos for authenticating captive portal, also make sure you understand these considerations:

If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.

DNS must return a response of 64KB or less to the hostname; otherwise, the AD connection test fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.

Port

The server's port.

Encryption

(Strongly recommended.) The encryption method to use:

- **`STARTTLS`**—encrypted LDAP connection
- **`LDAPS`**—encrypted LDAP connection
- **`None`**—unencrypted LDAP connection (unsecured traffic)

To communicate securely with an Active Directory server, see Connect securely to Active Directory or LDAP, on page 2328.

CA Certificate

The TLS/SSL certificate to use for authentication to the server. You must configure **`STARTTLS`** or **`LDAPS`** as the **Encryption** type to use a TLS/SSL certificate.

If you are using a certificate to authenticate, the name of the server in the certificate must match the server **Hostname / IP Address**. For example, if you use 10.10.10.250 as the IP address but **`computer1.example.com`** in the certificate, the connection fails.

Interface used to connect to Directory server

Required only for RA VPN authentication so the Secure Firewall Threat Defense can connect securely to your Active Directory server. This interface is not used for downloading users and groups, however.

You can choose only a routed interface group. For more information, see Interfaces, on page 2980.

Click one of these options:

- **Resolve via route lookup**: Use routing to connect to the Active Directory server.
- **Choose an interface**: Choose a specific managed device interface group to connect to the Active Directory server.

User synchronize fields

AD Primary Domain

For Microsoft Active Directory realms only. Domain for the Active Directory server where users should be authenticated.

Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different Microsoft AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

Enter query to look for users and groups

Base DN:

(Optional.) The directory tree on the server where the Cloud-Delivered Firewall Management Center should begin searching for user data.

Typically, the base distinguished name (DN) has a basic structure indicating the company domain name and operational unit. For example, the Security organization of the Example company might have a base DN of `ou=security,dc=example,dc=com`.

Group DN:

(Optional.) The directory tree on the server where the Cloud-Delivered Firewall Management Center should search for users with the group attribute. A list of supported group attributes is shown in Supported server object class and attribute names, on page 2279.

Note Neither the group name nor the organizational unit name can contain special characters like asterisk (*), equals (=), backslash (\) because users in those groups are not downloaded and cannot be used in identity policies.

Load Groups

Enables you to download users and groups for user awareness and user control.

Available Groups, Add to Include, Add to Exclude

Limits the groups that can be used in policy.

- Groups that are displayed in the **Available Groups** field are available for policy unless you move groups to the **Included Groups and Users** or **Excluded Groups and Users** field.
- If you move groups to the **Included Groups and Users** field, only those groups and users they contain are downloaded and user data is available for user awareness and user control.
- If you move groups to the **Excluded Groups and Users** field, all groups and users they contain *except* these are downloaded and available for user awareness and user control.
- To include users from groups that are not included, enter the user name in the field below **User Inclusion** and click **Add**.
- To exclude users from groups that are not excluded, enter the user name in the field below **User Exclusion** and click **Add**.

Note The users that are downloaded to the Cloud-Delivered Firewall Management Center is calculated using the formula **R = I - (E+e) + i**, where

- R is list of downloaded users
- I is included groups
- E is excluded groups
- e is excluded users
- i is included users

Synchronize Now
Click to synchronize groups and users with AD.

Begin automatic synchronization at
Enter the time and time interval at which to download users and groups from AD.

Create a proxy sequence

Create a proxy sequence to enable reliable user repository connectivity and failover support for managed devices.

A *proxy sequence* is one or more managed devices that can be used to communicate with an LDAP, Active Directory, or ISE/ISE-PIC server. It is necessary only if Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.)

Although you can use one managed device as a proxy sequence, we strongly recommend you set up two or more so that, in the event one managed device cannot communicate with Active Directory or ISE/ISE-PIC, another managed device can take over.

Note To use a proxy sequence with a SAML - Azure AD realm, the managed device must run version 7.4.2 or later.

Because of its limited capacity, we recommend you *do not* use a Secure Firewall 220 as a proxy sequence.

Before you begin

You must add at least two managed devices to Security Cloud Control, all of which must be able to communicate with Active Directory or ISE/ISE-PIC.

Follow these steps to create a proxy sequence:

Procedure

Step 1 Log in to the Security Cloud Control if you have not already done so.

Step 2 Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 3 Click **Add Proxy Sequence**.

Step 4 In the **Name** field, enter a name to identify the proxy sequence.

Step 5 (Optional.) In the **Description** field, enter a description for the proxy sequence.

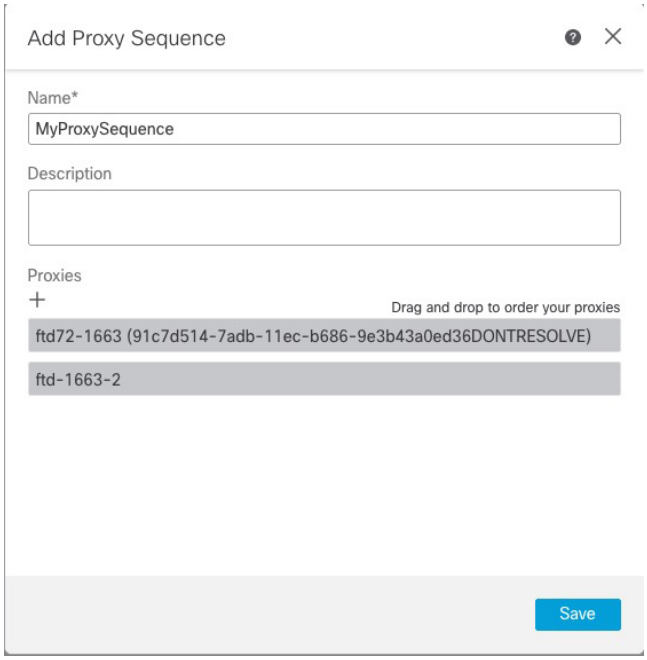
Step 6 Under Proxies, click **Add** (+).

Step 7 Click the name of each managed device to add to the sequence.

To narrow your search, enter all or part of a realm name in **Filter** field.

Click **OK**.

Step 8 In the Add Proxy Sequence dialog box, drag and drop the proxies in the order in which you want Security Cloud Control to search for them.
The following figure shows an example of a proxy sequence consisting of two proxies. The top proxy will be searched for users before the bottom proxy. Both proxies must be able to communicate with either Active Directory or ISE/ISE-PIC.



Step 9 Click **Save**.

What to do next

See Create an identity policy, on page 2507.

Create a Microsoft Azure AD (SAML) realm

You can use a Microsoft Azure Active Directory (AD) (now called *Entra ID*) realm for either passive authentication or active authentication.

Authentication options

Passive authentication

Passive authentications occur when a user authenticates with Cisco ISE.

You have these options, depending on your choice of user and group repository:

- To use Cisco ISE as a repository for users and to perform passive authentication using Entra ID. For more information, see:
  - How Entra ID and Cisco ISE authentication with resource owned password credentials works, on page 2291
  - How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS, on page 2292
- To download groups from Entra ID.

  For more information about setting up Entra ID, see Configure Microsoft Entra ID for passive authentication, on page 2293.

Active authentication

Active authentications occur when a user authenticates through preconfigured managed devices. Captive portal is another name for active authentication. Active authentication generally uses the same user repositories as passive authentication (the exceptions being ISE/ISE-PIC, and TS Agent, and the Passive Identity Agent, which are passive only).

To use Microsoft Entra ID as a captive portal requires users to authenticate with Entra ID. We refer to the realm as a Security Assertion Markup Language (SAML) realm because SAML is used to establish a trust relationship between:

- A *service provider* (the Secure Firewall Threat Defense device or devices to which authentication requests are sent).
- An *identity provider* (Microsoft Entra ID).

SAML is an open standard developed by the OASIS standards body; for more information, see the SAML Overview.

For more information, see Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal), on page 2302.

Create a Microsoft Azure AD realm for passive authentication

Create and configure a Microsoft Azure AD (SAML) (now called *Entra ID*) realm for passive authentication use with the Security Cloud Control.

This configuration enables passive authentication between Microsoft Entra ID and your firewall system through proper realm setup and access control policy integration.

Procedure

Step 1 Configure Microsoft Entra ID.

Several configuration tasks are required, including setting up an event hub, giving your application permission to the Microsoft Graph API, and enabling the audit log.

See Configure Microsoft Entra ID for passive authentication, on page 2293.

Step 2 Configure Cisco ISE.

The way you configure ISE depends on how users authenticate with your system. For more information, see Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD, on page 2293.

Step 3 Create a Cisco ISE identity source.

The identity source enables ISE to communicate with the Cloud-Delivered Firewall Management Center.

Step 4 Get the information required to configure your Microsoft Azure AD realm.

This information includes client and tenant IDs, client secret, and other information store in Microsoft Entra ID.

Step 5 Configure and verify your realm.

Test the realm's configuration before you start to use it in access control policies.

Create a Microsoft Azure AD (SAML) realm as discussed in Create a Microsoft Azure AD (SAML) realm, on page 2289

Step 6 Create access control policies and rules using your Microsoft Azure AD (SAML) realm.

Unlike other types of realms, you do not need to create an identity policy or associate the identity policy with an access control policy.

See Creating a basic access control policy, on page 2004 and Create and edit access control rules, on page 2066.

What to do next

See How Entra ID and Cisco ISE authentication with resource owned password credentials works, on page 2291.

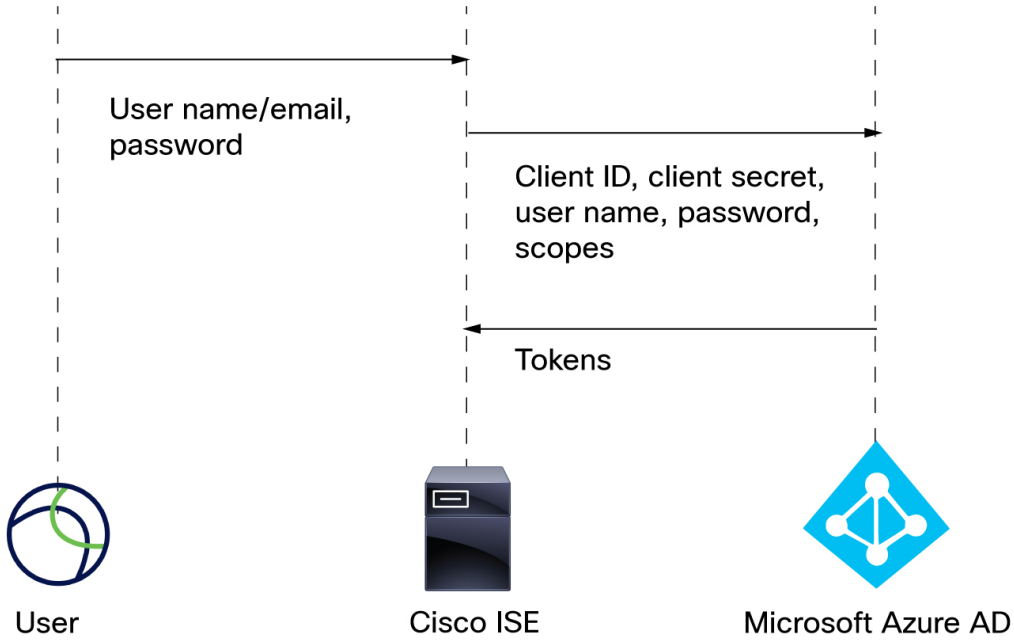
How Entra ID and Cisco ISE authentication with resource owned password credentials works

Summary

The key components involved in the Entra ID authentication with resource owned password credentials process are:

- User: Logs in with credentials using a VPN client
- Cisco Secure Client: VPN client that sends user credentials
- Entra ID (formerly Azure AD): Validates credentials and issues tokens
- Cisco ISE: Receives tokens and manages user sessions
- Cloud-Delivered Firewall Management Center: Receives user sessions from ISE

Workflow



These stages describe how Entra ID and Cisco ISE authentication with resource owned password credentials works:

1. The user logs in with a user name (or email address) and password using a VPN client like Cisco Secure Client.
2. The client ID, client secret, user name, password, and scopes are sent to Entra ID.
3. Tokens are sent from Entra ID to Cisco ISE, which sends user sessions to the Cloud-Delivered Firewall Management Center.

   For details about configuring Cisco ISE, see Configure ISE 3.0 REST ID with Azure Active Directory.

   Additional resource: Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials on learn.microsoft.com.

How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS

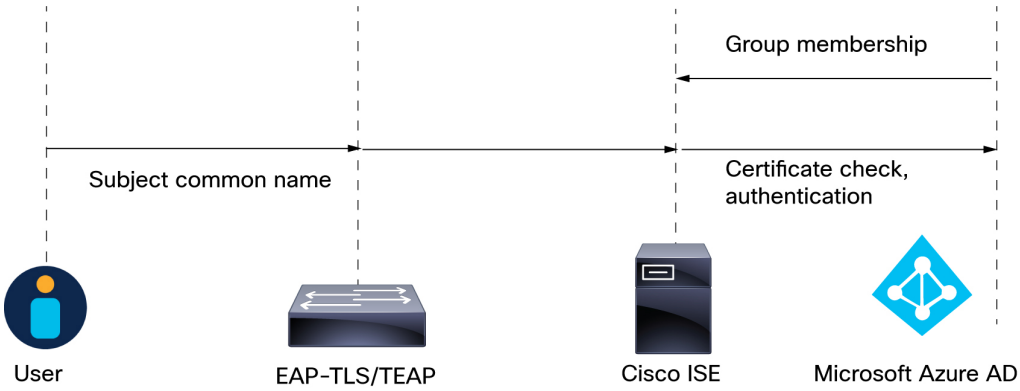
Tunnel Extensible Authentication Protocol (TEAP), defined by RFC7170, can be used with ISE and the Security Cloud Control for secure authentication.

Summary

The key components involved in this authentication process are:

- User device: Provides certificate through EAP-TLS or TEAP protocol
- ISE: Validates certificate and performs user lookup through Azure Graph API
- Entra ID: Provides user attributes and group information for authorization

Workflow



These stages describe how the authentication process works and are based on Configure Cisco ISE 3.2 EAP-TLS with Microsoft Azure Active Directory:

1. The user's certificate is sent to ISE through EAP-TLS or TEAP with EAP-TLS as the inner method.
2. ISE evaluates the user's certificate (validity period, trusted certificate authority, certificate revocation list, and so on).
3. ISE takes the certificate subject name (CN) and performs a look-up to the Azure Graph API to fetch the user's groups and other attributes. This is referred to by Azure as User Principal name (UPN).
4. ISE authorization policies are evaluated against the user's attributes returned from Azure.

Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD

In a Microsoft Azure AD (SAML)Microsoft Azure AD (now called *Entra ID*) realm, it's the responsibility of ISE to send user sessions (login, logout) to the Cloud-Delivered Firewall Management Center. This topic discusses how to set up ISE for use with an Azure AD realm.

Resource owned password credentials authentication

To use ISE with Microsoft Azure AD (SAML)Microsoft Azure ADMicrosoft Azure AD implemented through Representational State Transfer (REST) Identity (ID) service with the help of Resource Owner Password Credentials (ROPC), see Configure ISE 3.0 REST ID with Azure Active Directory.

TEAP/EAP-TLS

To use ISE with authorization policies based on Azure AD group membership and other user attributes with EAP-TLS or TEAP as the authentication protocols, see Configure Cisco ISE 3.2 EAP-TLS with Microsoft Azure Active Directory.

What to do next

Get required information For Your Microsoft Azure AD realm, on page 2296

Configure Microsoft Entra ID for passive authentication

Configure Microsoft Entra ID to enable passive authentication and user group synchronization with the Security Cloud Control.

This task provides basic information about how to set up a Microsoft Entra ID (formerly called *Microsoft Azure Active Directory (AD)*) as a realm you can use with the Security Cloud Control. We expect you to already be familiar with Entra ID; if not, consult documentation or a support resource before you get started.

Procedure

Step 1 Grant your Entra ID application the Microsoft Graph permissions.

Grant your Entra ID application the following permissions to Microsoft Graph as discussed in Authorization and the Microsoft Graph Security API on the Microsoft site:

- Reader role
- User.Read.All permission
- Group.Read.All permission

This permission enables the Security Cloud Control to download users and groups from Entra ID the first time.

Required information from this step for setting up the Entra ID realm in the Security Cloud Control:

- Name of the app you registered
- **Application (client) ID**
- **Client secret**
- **Directory (tenant) ID**

Step 2 Set up an event hub.

Set up the event hub as discussed in Quickstart: Create an event hub using Azure portal on the Microsoft site. The Security Cloud Control uses the event hub audit log to download periodic updates to users and groups.

More information: Features and terminology in Azure Event Hubs.

Important
You must choose the **Standard** pricing tier or better. If you choose **Basic**, the realm cannot be used.

Required information from this step for setting up the Entra ID realm in the Security Cloud Control:

- Namespace **Name**
- **Connection string—primary key**
- Event Hub **Name**
- Consumer group **Name**

Step 3 Enable the audit log.

Enable the audit log as discussed in Tutorial: Stream Azure Active Directory logs to an Azure event hub on the Microsoft site.

Step 4 Configure Cisco ISE for Entra ID.

To send user session information to the Security Cloud Control, configure Cisco ISE for Entra ID as discussed in Configure ISE 3.0 REST ID with Azure Active Directory.

What to do next

See Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD, on page 2293.

Configure Entra ID basic settings

Set up Entra ID realm authentication to enable the Cloud-Delivered Firewall Management Center to download and periodically update users and groups from Entra ID.

This configuration enables SAML realm authentication for Entra ID (formerly Azure AD) in Secure Firewall Management Center, allowing integration with Microsoft Graph services and audit logging capabilities.

Before you begin

Follow these steps to configure Entra ID basic settings:

Procedure

Step 1 Give your application the Microsoft Graph permission.

Grant your Entra ID application (formerly called *Azure AD*) the following permissions to Microsoft Graph as discussed in Authorization and the Microsoft Graph Security API on the Microsoft site:

- Reader role
- User.Read.All permission
- Group.Read.All permission

This permission enables the Cloud-Delivered Firewall Management Center to download users and groups from Entra ID the first time.

Required information from this step for setting up the Entra ID realm in the Cloud-Delivered Firewall Management Center:

- Name of the app you registered
- **Application (client) ID**
- **Client secret**
- **Directory (tenant) ID**

Step 2 Set up an event hub.

Set up the event hub as discussed in Quickstart: Create an event hub using Azure portal on the Microsoft site. The Cloud-Delivered Firewall Management Center uses the event hub audit log to download periodic updates to users and groups.

More information: Features and terminology in Azure Event Hubs.

Important

You must choose the **Standard** pricing tier or better. If you choose **Basic**, the realm cannot be used.

Required information from this step for setting up the Entra ID realm in the Security Cloud Control:

- Namespace **Name**
- **Connection string—primary key**
- Event Hub **Name**
- Consumer group **Name**

Step 3 Enable the audit log.

Enable the audit log as discussed in Tutorial: Stream Azure Active Directory logs to an Azure event hub on the Microsoft site.

Get required information For Your Microsoft Azure AD realm

This task enables you to collect the necessary credentials and configuration information from Microsoft Entra ID (formerly Azure AD) and Event Hubs to establish a SAML authentication realm in your Secure Firewall Management Center.

This task explains how to get the information required to set up a Microsoft Azure AD (now called *Entra ID*) realm in the Security Cloud Control. You might have already obtained this information when you set up Microsoft Entra ID as discussed in Configure Microsoft Entra ID for passive authentication, on page 2293.

Procedure

Step 1 Log in to https://portal.azure.com/ as a user with at least the Product Designer role.

Step 2 At the top of the page, click **Microsoft Entra ID**.

Step 3 In the left column, click **App Registrations**.

a) If necessary, filter the list of displayed apps to show the one you want to use.

b) Click the name of your app.



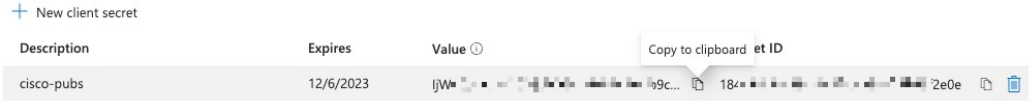
c) Click **Copy** ( ) next to the following values on this page and paste those values to a text file.

- **Application (Client) ID**
- **Directory (tenant) ID**

d) Click **Client Credentials**.

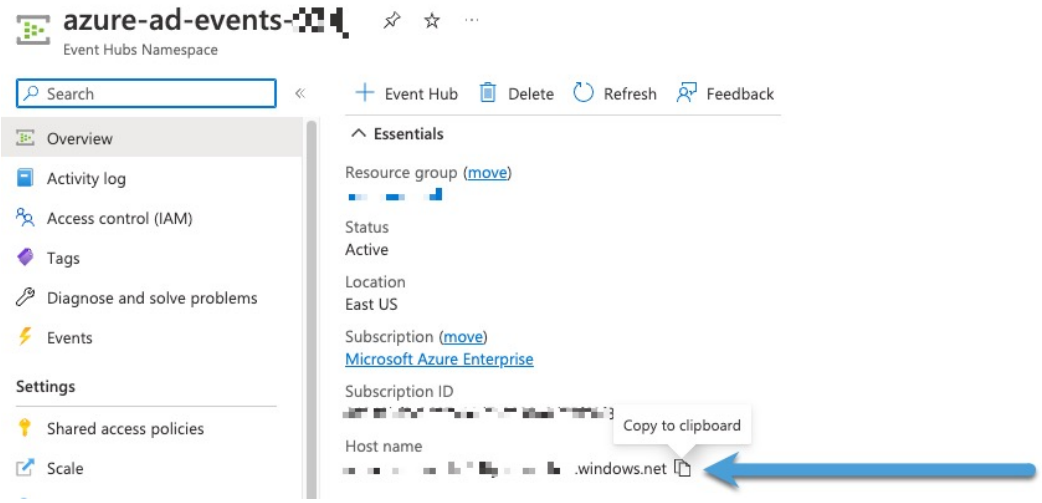
e) Unless you already know the client secret *value* (as opposed to the client *secret ID*), you must create a new client secret as follows:

1. Click **New Client Secret**.
2. Enter the required information in the provided fields.
3. Click **Add**.
4. Click **Copy** ( ) next to Value as the following figure shows.

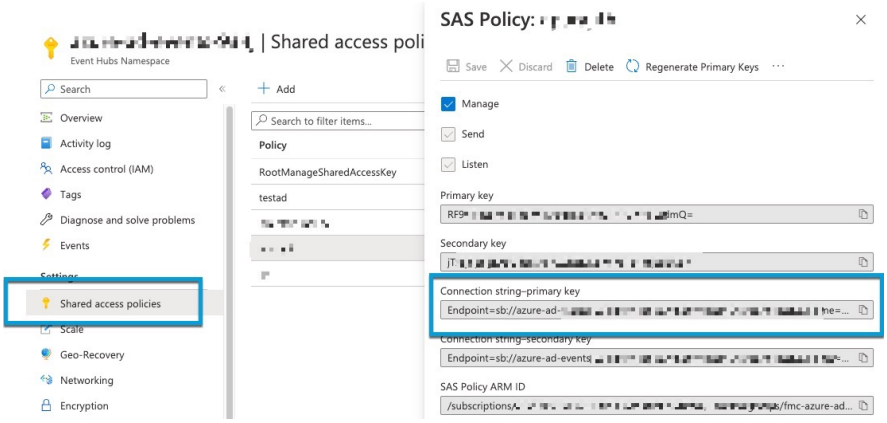


Step 4 From https://portal.azure.com/, click **Event Hubs** > **(name of an event hub)**.

a) In the right pane, click **Copy** ( ) next to the value of **Host name** and paste the value to the clipboard. This is your event hub host name.

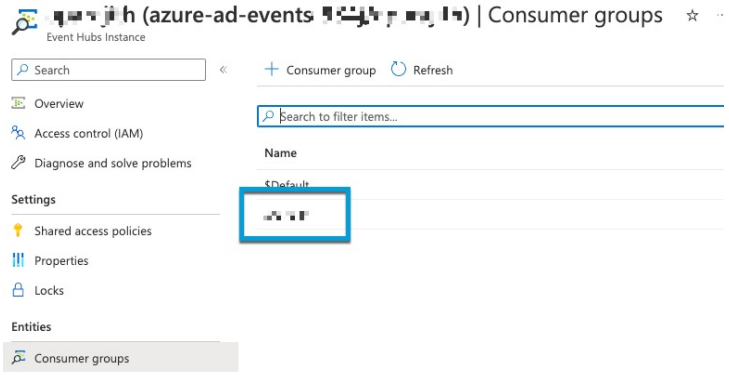


b) Write down or copy to a text file the name of the event hub (same as the **Event Hubs Namespace** at the top of the page).
c) In the left pane, under Settings, click **Shared access policies**.
d) Click the name of a policy.
e) Click **Copy** ( ) next to **Connection string-primary key**.



Step 5 Click **Overview** > **Entities** > **Event Hubs** > **(name of an event hub)** > **Entities** > **Consumer Groups**.

Write down this value or copy it to the clipboard. This is your consumer group name.



Step 6 In the left pane, click **Overview**.

Step 7 Click **Copy** ( ) next to **Namespace**.



This is your event hubs topic name.

Create a Microsoft Azure AD (SAML) realm for passive authentication

A Microsoft Azure AD (SAML) realm for passive authentication retrieves groups from Microsoft Azure Active Directory (now called *Entra ID*) and logged-in user session data from Cisco ISE to authenticate users.

Authentication options

You have the following authentication options:

- Resource owned password credentials (ROPC): Enables users to log in with a client like Cisco Secure Client using a user name and password. ISE sends user sessions to the Cloud-Delivered Firewall Management Center. For more information, see How Entra ID and Cisco ISE authentication with resource owned password credentials works, on page 2291.

  Additional resource: Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials on learn.microsoft.com.

- Extensible Authentication Protocol (EAP) Chaining with Tunnel-based Extensible Authentication Protocol (TEAP) and Transport Layer Security (TLS), abbreviated EAP/TEAP-TLS: TEAP is a tunnel-based EAP method that establishes a secure tunnel and executes other EAP methods under the protection of that secured tunnel. ISE is used to validate user credentials and to send user sessions to the Cloud-Delivered Firewall Management Center. For more information, see How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS, on page 2292.

To configure the realm, complete all tasks in this order:

1. Configure Entra ID basic settings, on page 2295.
2. Get required information for your realm as discussed in Get required information For Your Microsoft Azure AD realm, on page 2296.
3. Microsoft Azure AD (SAML) realm: SAML details, on page 2299.

Microsoft Azure AD (SAML) realm: SAML details

Set up a Microsoft Entra ID (SAML) realm to enable user authentication through passive or active methods with integration between the Secure Firewall Management Center and Azure AD and identity providers.

This task discusses the first step in a multi-step wizard that creates a Microsoft Azure AD (SAML) realm. You must complete all steps in the wizard to set up your realm. The steps are different depending on whether you create the realm for active or passive authentication.

Before you begin

Complete all tasks before you create your realm:

- (Passive authentication with Cisco ISE only.) If you're using Cisco ISE as the repository for users and groups, set up ISE:
  - How Entra ID and Cisco ISE authentication with resource owned password credentials works, on page 2291
  - How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS, on page 2292
- To use Entra ID as the repository for users and groups, see Configure Microsoft Entra ID for passive authentication, on page 2293.
- Get required information for your realm as discussed in Get required information For Your Microsoft Azure AD realm, on page 2296.

Follow these steps to configure a Microsoft Entra ID (SAML) realm:

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 3 Click **Add Realm** > **SAML - Azure AD**.

Step 4 Enter the required information.

Item	Description
Name	Unique name to identify the realm.
Description	(Optional.) Description of the realm.
Identity Provider	Always displays **Azure AD**.
Configuration Type	Click one of these options: • **Passive Authentication with ISE** for passive authentication. • **Passive authentication or captive portal with Azure AD** for to use Entra ID as a user store for either passive authentication or active authentication (that is, captive portal).

Step 5 Click **Next**.

The realm creation wizard advances to the next step based on your selected configuration type.

What to do next

Continue with one of these tasks:

- Passive authentication: Microsoft Azure AD (SAML) realm: Azure AD details, on page 2300.
- Active authentication: Microsoft Azure AD (SAML) realm: SAML service provider (SP) metadata, on page 2312.

Microsoft Azure AD (SAML) realm: Azure AD details

Configure the Entra ID (SAML) realm to enable authentication and user data management through Azure AD integration.

This task discusses one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

Procedure

Step 1 Continue from the preceding step in the wizard.

Step 2 Enter the following information.

Item	Description
Name	Enter a unique name to identify this realm.
Client Secret **Tenant ID** **Event Hubs Host Name** **Event Hub Name** **Event Hub Connection String**	Enter the information you found as discussed in:. • Passive authentication: Create a Microsoft Azure AD realm for passive authentication, on page 2290 • Active authentication: Get required information for your Microsoft Azure AD realm (active authentication only), on page 2307
(Optional.) **User Groups**	Slide to **Slider enabled** ( ) to specify groups to include or exclude from policy.
(Optional.) **Excluded User Groups**	If you enter one or more group names in this field, all groups and users they contain *except* these are downloaded and available for user awareness and user control. Enter one group name per line followed by a line break. Group names are case-sensitive.
(Optional.) **Included User Groups**	If you enter one or more group names in this field, only those groups and users they contain are downloaded and user data is available for user awareness and user control. Enter one group name per line followed by a line break. Group names are case-sensitive. We recommend that you use the Included User Groups field to download only those groups that you intend to monitor with Secure Firewall.

Step 3 Click **Test**.
Make sure the test connection succeeds before you continue to the next step.

Step 4 Click **Next**.

The Entra ID (SAML) realm configuration is completed with Azure AD details, and the connection test confirms successful integration.

Configure Microsoft Azure AD (SAML) realm user session timeout

Set the session timeout to define how long inactive ISE or captive portal user sessions remain active before automatic termination in a Microsoft Entra ID (SAML) realm configuration.

This task discussed one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

This option sets the number of seconds before an inactive session is terminated by the system.

Procedure

Step 1 Continue from the preceding step in the wizard.

Step 2 Enter the following information.

Item	Description
ISE users	Default is 1440 minutes (24 hours).
Captive portal users	Default is 1440 minutes (24 hours).

After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Cloud-Delivered Firewall Management Center as Unknown (except for **Failed Captive Portal Users**).

Step 3 Click **Save**.

Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal)

Create and configure a Microsoft Azure Active Directory (Entra ID) realm to enable secure user authentication, certificate management, and access control for protected network resources.

This topic discusses the high-level tasks of creating a Microsoft Azure Active Directory (AD) realm (now called *Entra ID*) for use with the Security Cloud Control.

Before you begin

If you enabled change management, you must open or edit, assign, and approve a ticket for each of the following objects before you can create the realm:

- Base URL
- Service provider certificate enrollment (PKCS12 format)
- Identity provider certificate enrollment (manual format)
- The realm itself (create and assign the ticket until realm creation is complete, then approve it)

For more information, see Opening a ticket for configuration changes, on page 176 and Policies and objects that support change management, on page 169.

Follow these steps to create a Microsoft Azure Active Directory (Entra ID) realm:

Procedure

Step 1 Create a fully-qualified host name (FQDN) using your DNS server and upload the Firewall Threat Defense's internal certificate to the Security Cloud Control.

You can consult a resource such as this one if you've never done it before. Specify the IP address of a routed interface on one of the devices managed by your Security Cloud Control.

Consult a DNS server reference.

Step 2 Create a network object with an associated internal certificate.

See Create a network object, on page 2985.

Step 3 Get a signed certificate and upload it to the Secure Firewall Threat Defense to which Entra ID authentication requests will be sent.

The certificate should be signed by a trusted Certificate Authority (CA) and delivered to you in **.p12** format (also referred to as PKCS#12; see also this article on ssl.com).

For background, see the section on public key infrastructure in Cisco Secure Firewall Management Center Device Configuration Guide or stackoverflow.com.

To upload the signed certificate, see Install a certificate using a PKCS12 file, on page 3087.

Step 4 Configure Microsoft Entra ID basic settings.

Several configuration tasks are required, including setting up an event hub, giving your application permission to the Microsoft Graph API, and enabling the audit log.

See Configure Entra ID basic settings, on page 2295.

Step 5 Create a single sign-on (SSO) app in Entra ID.

The SSO app enables users that request access to a protected network resource to authenticate with Entra ID. The SSO app has both the federation XML that you can use to simplify realm creation as well as the identity provider certificate the Secure Firewall Threat Defense requires to security authenticate with Entra ID.

See Configure a single sign-on (SSO) app in Entra ID, on page 2305.

Step 6 Get the information required to configure your Microsoft Azure AD (SAML) realm.

This information includes client and tenant IDs, client secret, and other information store in Microsoft Entra ID.

See Get required information For Your Microsoft Azure AD realm, on page 2296.

Step 7 Configure a decryption policy with a **Decrypt - Resign** rule for the Azure Authentication Service so users can access web pages using the HTTPS protocol.

The Microsoft Azure AD (SAML) realm can authenticate users only if the HTTPS traffic is decrypted before the traffic is sent to the realm. The Microsoft Azure AD (SAML) realm itself is seen by the system as the **Azure Authentication Service** application.

Create a decryption rule with decrypt - resign action, on page 2307.

Step 8 Do the following:

a) Create identity policy with an active authentication rule.

The identity policy enables selected users in your realm access resources after authenticating with the Microsoft Azure AD (SAML) realm.

For more information, see Create an identity policy, on page 2507.

b) Create access control policies and rules using your Microsoft Azure AD realm.

Unlike other types of realms, you do not need to create an identity policy or associate the identity policy with an access control policy.

See Creating a basic access control policy, on page 2004 and Create and edit access control rules, on page 2066.

c) Associate the identity and decryption policies with the access control policy from step 3.

This final step enables the system to authenticate users with the Microsoft Azure AD (SAML) realm.

For more information, see Associate prefilter, decryption, and identity policies with an access control, on page 2016.

What to do next

See Configure Entra ID basic settings, on page 2295.

Configure Entra ID basic settings

Set up Entra ID realm authentication to enable the Cloud-Delivered Firewall Management Center to download and periodically update users and groups from Entra ID.

This configuration enables SAML realm authentication for Entra ID (formerly Azure AD) in Secure Firewall Management Center, allowing integration with Microsoft Graph services and audit logging capabilities.

Before you begin

Follow these steps to configure Entra ID basic settings:

Procedure

Step 1 Give your application the Microsoft Graph permission.

Grant your Entra ID application (formerly called *Azure AD*) the following permissions to Microsoft Graph as discussed in Authorization and the Microsoft Graph Security API on the Microsoft site:

- Reader role
- User.Read.All permission
- Group.Read.All permission

This permission enables the Cloud-Delivered Firewall Management Center to download users and groups from Entra ID the first time.

Required information from this step for setting up the Entra ID realm in the Cloud-Delivered Firewall Management Center:

- Name of the app you registered
- **Application (client) ID**
- **Client secret**
- **Directory (tenant) ID**

Step 2 Set up an event hub.

Set up the event hub as discussed in Quickstart: Create an event hub using Azure portal on the Microsoft site. The Cloud-Delivered Firewall Management Center uses the event hub audit log to download periodic updates to users and groups.

More information: Features and terminology in Azure Event Hubs.

Important
You must choose the **Standard** pricing tier or better. If you choose **Basic**, the realm cannot be used.

Required information from this step for setting up the Entra ID realm in the Security Cloud Control:

- Namespace **Name**
- **Connection string—primary key**
- Event Hub **Name**
- Consumer group **Name**

Step 3 Enable the audit log.

Enable the audit log as discussed in Tutorial: Stream Azure Active Directory logs to an Azure event hub on the Microsoft site.

Configure a single sign-on (SSO) app in Entra ID

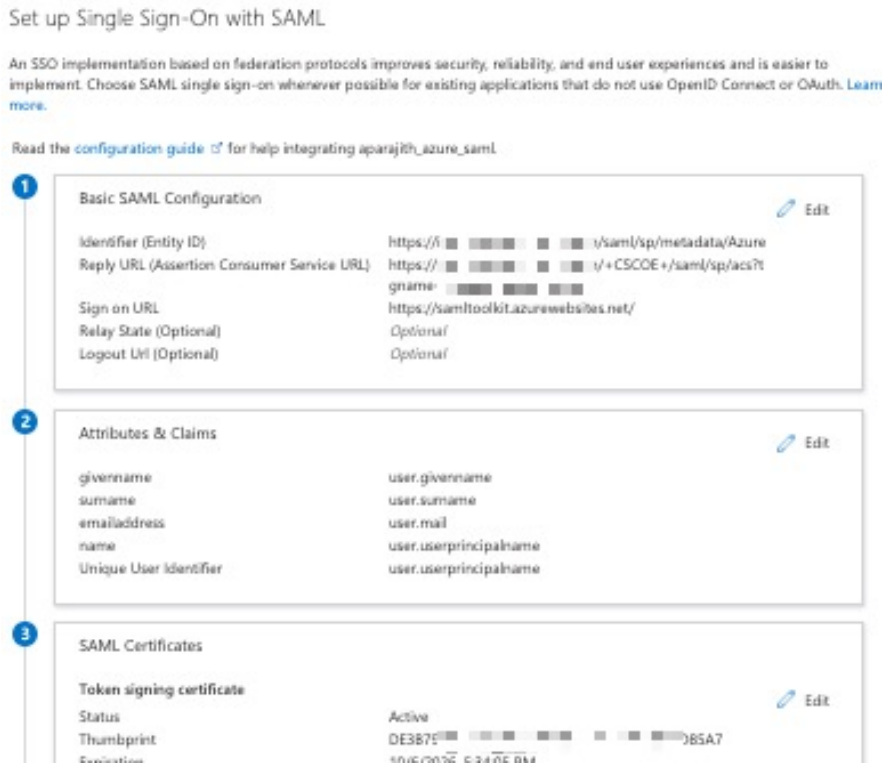
Create a Microsoft Entra ID app to enable single sign-on (SSO) for users accessing a protected network.

This task discusses how to create an app in Microsoft Entra ID to handle single sign-on (SSO) from Entra ID when a network user attempts to access a protected network.

Procedure

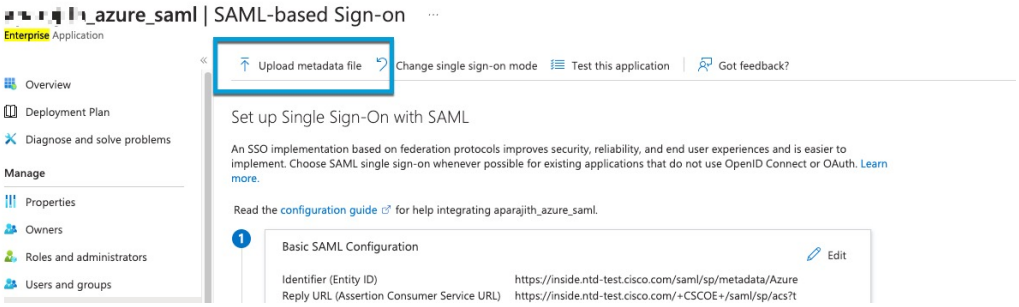
Step 1 In the Microsoft Entra ID portal, click **Enterprise Applications** on the home page and follow the instructions in Configure Microsoft Entra SSO on learn.microsoft.com.

The following figure shows part of the SSO app configuration. You must provide some of the information displayed on this page when you configure the Microsoft Azure AD (SAML) realm. For more information, see Get required information for your Microsoft Azure AD realm (active authentication only), on page 2307.



Step 2 (Optional.) If you have already configured the Microsoft Azure AD (SAML) realm, click **Upload metadata file** at the top of the page to quickly provide configuration values for the SSO app.

The following figure shows an example.



Step 3 Add users and groups to your app as discussed in Add a user account to an enterprise application on learn.microsoft.com

What to do next

See Get required information for your Microsoft Azure AD realm (active authentication only), on page 2307.

Create a decryption rule with decrypt - resign action

This task creates a decryption policy that enables secure inspection and authentication of encrypted traffic using the firewall and the Azure Authentication Service.

This part of the procedure discusses how to create a decryption policy to decrypt and resign traffic before the traffic reaches the SAML realm. The realm can authenticate traffic only after it has been decrypted.

Procedure

Step 1 If you haven't done so already, log in to the Security Cloud Control.

Step 2 If you haven't done so already, create an internal certificate authority object to decrypt TLS/SSL traffic as discussed in PKI objects, on page 2986.

Step 3 Click **Policies** > **FTD Policies**.

Step 4 Click **Policies** > **Access Control heading** > **Decryption**.

Step 5 Click **New Policy**.

Step 6 Enter a **Name** and choose a **Default Action** for the policy. Default actions are discussed in Rule-based decryption policy default actions, on page 2179.

Step 7 Click **Save**.

Step 8 At the top of the page, click **Save**.

Get required information for your Microsoft Azure AD realm (active authentication only)

This task explains how to get the information required to set up a Microsoft Azure AD realm (now called *Entra ID*) in the Cloud-Delivered Firewall Management Center.

Follow these steps to get required information for your Microsoft Azure AD realm:

Procedure

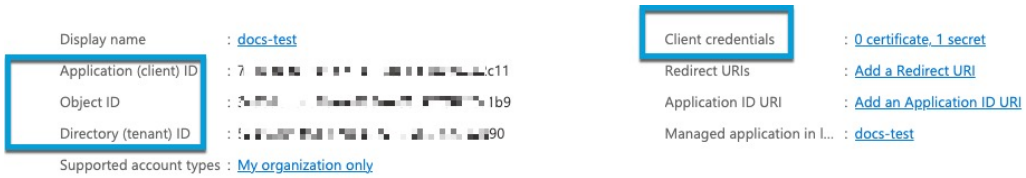
Step 1 Log in to https://portal.azure.com/ as a user with at least the Product Designer role.

a) At the top of the page, click **Microsoft Entra ID**.

b) In the left column, click **App Registrations**.

c) If necessary, filter the list of displayed apps to show the one you want to use.

d) Click the name of your app.



e) Click **Copy** ( ) next to the following values on this page and paste those values to a text file.

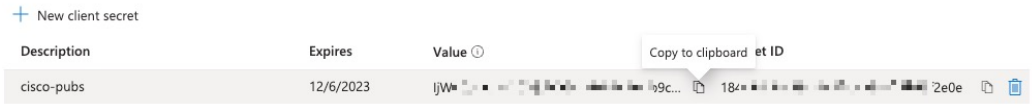
- **Application (Client) ID**

- **Directory (tenant) ID**

f) Click **Client Credentials**.

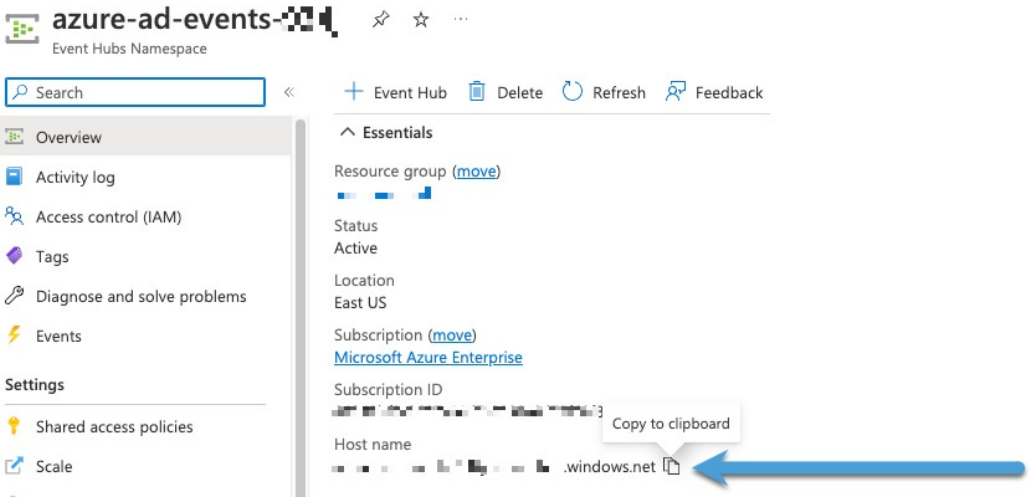
Step 2 Unless you already know the client secret *value* (as opposed to the client *secret ID*), you must create a new client secret as follows:

a) Click **New Client Secret**.
b) Enter the required information in the provided fields.
c) Click **Add**.
d) Click **Copy** ( ) next to Value as the following figure shows.

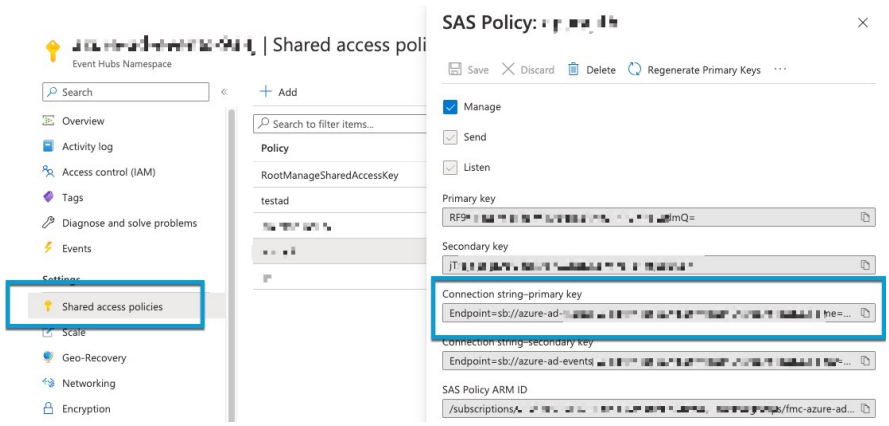


Step 3 From https://portal.azure.com/, click **Event Hubs** > **(name of an event hub)**.

a) In the right pane, click **Copy** ( ) next to the value of **Host name** and paste the value to the clipboard. This is your event hub host name.

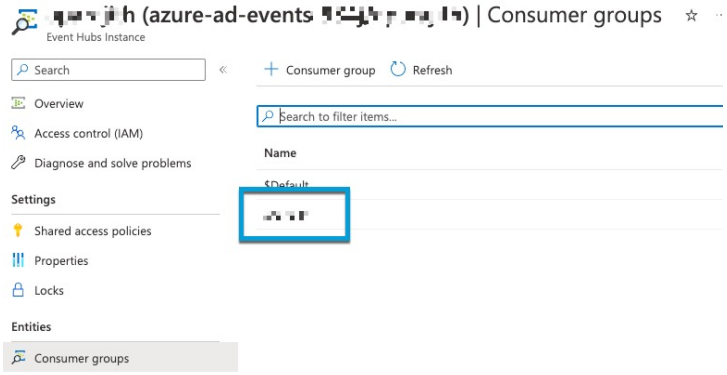


b) Write down or copy to a text file the name of the event hub (same as the **Event Hubs Namespace** at the top of the page).
c) In the left pane, under Settings, click **Shared access policies**.
d) Click the name of a policy.
e) Click **Copy** ( ) next to **Connection string-primary key**.



Step 4 Click **Overview** > **Entities** > **Event Hubs** > **(name of an event hub)** > **Entities** > **Consumer Groups**.

Write down the following value or copy it to the clipboard. This is your consumer group name.



Step 5 In the left pane, click **Overview**.

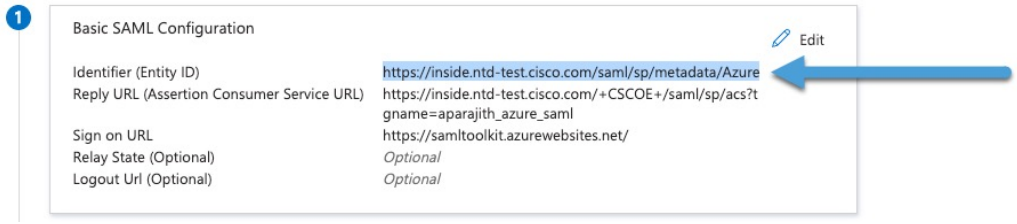
Step 6 Click **Copy** ( ) next to **Namespace**.



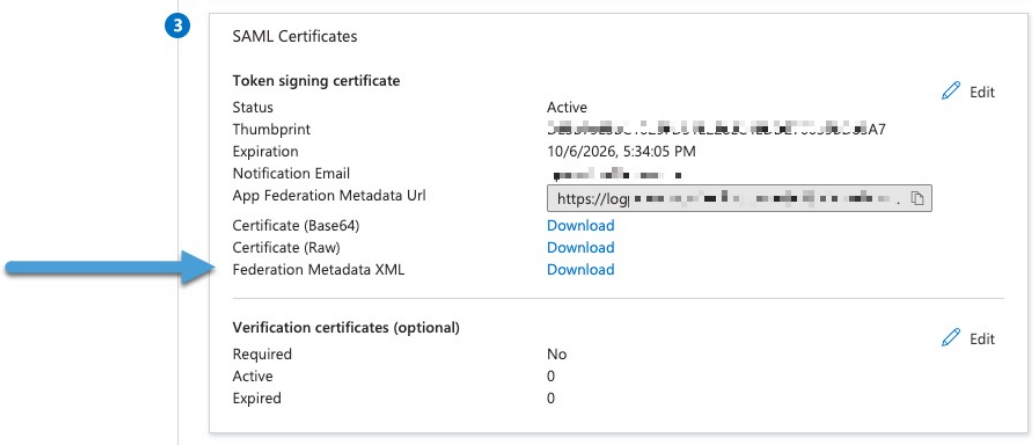
This is your event hubs topic name.

Step 7 Return to the home page and log in if necessary: https://portal.azure.com/#home.

a) Click **Microsoft Entra ID**.
b) In the left pane, click **Enterprise Applications.**
c) If necessary, filter the list of applications to locate yours.
d) Click the name of your enterprise application.
e) Click **Get Started** under Set up single sign on.
f) On your SSO app page, copy the value of Identifier (Entity ID) to the clipboard.
The following figure shows an example.

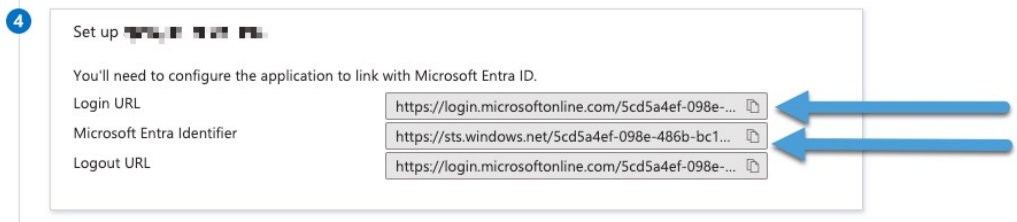


g) On your SSO app page, click the **Download** link next to Federation Metadata XML, as this figure shows. The following figure shows an example.



h) If you've already set up your SSO app, *you can stop here*. The Federation Metadata XML contains all the information required to configure the identity provider in the Cloud-Delivered Firewall Management Center.

i) (Optional if you already downloaded the federation XML.) Click **Copy** ( ) next to both of these values and save them to a text file.
This figure shows an example.



What to do next

See Create a decryption rule with decrypt - resign action, on page 2307.

Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal)

Create a Microsoft Entra ID SAML realm for active authentication where users must authenticate with Microsoft Entra ID before accessing protected resources.

In active authentication (also referred to as *captive portal*), Microsoft Entra ID is the user store. When a user tries to access a protected resource as defined in access control rules, the user must authenticate with Microsoft Entra ID first.

To configure the realm, complete all tasks in the following order.

Procedure

Step 1 Configure Entra ID basic settings, on page 2295.

Step 2 Configure a single sign-on (SSO) app in Entra ID, on page 2305.

Step 3 Get required information for your Microsoft Azure AD realm (active authentication only), on page 2307

Microsoft Azure AD (SAML) realm: SAML details

Set up a Microsoft Entra ID (SAML) realm to enable user authentication through passive or active methods with integration between the Secure Firewall Management Center and Azure AD and identity providers.

This task discusses the first step in a multi-step wizard that creates a Microsoft Azure AD (SAML) realm. You must complete all steps in the wizard to set up your realm. The steps are different depending on whether you create the realm for active or passive authentication.

Before you begin

Complete all tasks before you create your realm:

- (Passive authentication with Cisco ISE only.) If you're using Cisco ISE as the repository for users and groups, set up ISE:
  - How Entra ID and Cisco ISE authentication with resource owned password credentials works, on page 2291
  - How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS, on page 2292
- To use Entra ID as the repository for users and groups, see Configure Microsoft Entra ID for passive authentication, on page 2293.
- Get required information for your realm as discussed in Get required information For Your Microsoft Azure AD realm, on page 2296.

Follow these steps to configure a Microsoft Entra ID (SAML) realm:

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 3 Click **Add Realm** > **SAML - Azure AD**.

Step 4 Enter the required information.

Item	Description
Name	Unique name to identify the realm.
Description	(Optional.) Description of the realm.
Identity Provider	Always displays **Azure AD**.
Configuration Type	Click one of these options: • **Passive Authentication with ISE** for passive authentication. • **Passive authentication or captive portal with Azure AD** for to use Entra ID as a user store for either passive authentication or active authentication (that is, captive portal).

Step 5 Click **Next**.

The realm creation wizard advances to the next step based on your selected configuration type.

What to do next

Continue with one of these tasks:

- Passive authentication: Microsoft Azure AD (SAML) realm: Azure AD details, on page 2300.
- Active authentication: Microsoft Azure AD (SAML) realm: SAML service provider (SP) metadata, on page 2312.

Microsoft Azure AD (SAML) realm: SAML service provider (SP) metadata

Configure the SAML service provider metadata to enable secure authentication between your Secure Firewall Threat Defense and Microsoft Entra ID.

This task discusses one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

Before you begin

Complete the tasks discussed in Microsoft Azure AD (SAML) realm: SAML details, on page 2299

Follow these steps to configure the SAML service provider metadata:

Procedure

Step 1 Continue from Microsoft Azure AD (SAML) realm: SAML details, on page 2299.

Step 2 Enter the following information.

Item	Description
Base URL	From the list, click the network object you previously created. Network users are directed to this URL when they try to access protected network resources. You can also click **Add** (+) to create an object now.
Entity ID	Your SSO app's entity ID.
Assertion Consumer Services (ACS) URL	Automatically generated from the preceding values.
Service Provider Certificate	From the list, click the certificate to use to decrypt requests to the Secure Firewall Threat Defense. You can also click **Add** (+) to create an object now.
Download Service Provider Metadata	(Optional.) Download the metadata associated with the service provider (that is, managed device) to simplify configuring your Microsoft Entra ID SSO app.

Step 3 Click **Next**.

What to do next

Microsoft Azure AD (SAML) Realm: SAML identity provider (IdP) metadata, on page 2313.

Microsoft Azure AD (SAML) Realm: SAML identity provider (IdP) metadata

Configure SAML identity provider metadata to enable authentication with Microsoft Entra ID in a secure firewall environment.

This task discusses one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

Before you begin

Complete the tasks discussed in Microsoft Azure AD (SAML) realm: SAML service provider (SP) metadata, on page 2312.

Follow these steps to configure SAML identity provider metadata:

Procedure

Step 1 Continue from Microsoft Azure AD (SAML) realm: SAML service provider (SP) metadata, on page 2312.

Step 2 If you previously downloaded the Entra ID SSO app's federation XML, click **Upload XML** and upload it now. You can then skip the next step.

Step 3 Enter the following information.

Item	Description
Entity ID	Enter your identity provider's entity ID.
Single Sign on (SSO) URL	Enter the app's SSO URL.
IdP Certificate	From the list, click the certificate to use to authenticate with Microsoft Entra ID. You can also click **Add** (+) to create an object now.

Step 4 Click **Next**.

What to do next

Microsoft Azure AD (SAML) realm: SAML details, on page 2299.

Microsoft Azure AD (SAML) realm: Azure AD details

Configure the Entra ID (SAML) realm to enable authentication and user data management through Azure AD integration.

This task discusses one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

Procedure

Step 1 Continue from the preceding step in the wizard.

Step 2 Enter the following information.

Item	Description
Name	Enter a unique name to identify this realm.
Client Secret	Enter the information you found as discussed in:. • Passive authentication: Create a Microsoft Azure AD realm for passive authentication, on page 2290 • Active authentication: Get required information for your Microsoft Azure AD realm (active authentication only), on page 2307
Tenant ID	
Event Hubs Host Name	
Event Hub Name	
Event Hub Connection String	
(Optional.) **User Groups**	Slide to **Slider enabled** (toggle) to specify groups to include or exclude from policy.

Item	Description
(Optional.) **Excluded User Groups**	If you enter one or more group names in this field, all groups and users they contain *except* these are downloaded and available for user awareness and user control. Enter one group name per line followed by a line break. Group names are case-sensitive.
(Optional.) **Included User Groups**	If you enter one or more group names in this field, only those groups and users they contain are downloaded and user data is available for user awareness and user control. Enter one group name per line followed by a line break. Group names are case-sensitive. We recommend that you use the Included User Groups field to download only those groups that you intend to monitor with Secure Firewall.

Step 3 Click **Test**.
Make sure the test connection succeeds before you continue to the next step.

Step 4 Click **Next**.

The Entra ID (SAML) realm configuration is completed with Azure AD details, and the connection test confirms successful integration.

Configure Microsoft Azure AD (SAML) realm user session timeout

Set the session timeout to define how long inactive ISE or captive portal user sessions remain active before automatic termination in a Microsoft Entra ID (SAML) realm configuration.

This task discussed one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

This option sets the number of seconds before an inactive session is terminated by the system.

Procedure

Step 1 Continue from the preceding step in the wizard.

Step 2 Enter the following information.

Item	Description
ISE users	Default is 1440 minutes (24 hours).
Captive portal users	Default is 1440 minutes (24 hours).

After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Cloud-Delivered Firewall Management Center as Unknown (except for **Failed Captive Portal Users**).

Step 3 Click **Save**.

Create a Microsoft Azure AD Realm

You can use a Microsoft Azure Active Directory (AD) realm with Cisco ISE to authenticate users and get user sessions for user control. We get groups from Azure AD and logged-in user session data from Cisco ISE.

Authentication options

You have two authentication options:

- Resource owned password credentials (ROPC): Enables users to log in with a client like AnyConnect using a user name and password. ISE sends user sessions to the Cloud-Delivered Firewall Management Center. For more information, see How Entra ID and Cisco ISE authentication with resource owned password credentials works, on page 2291.

  Additional resource: Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials on learn.microsoft.com.

- Extensible Authentication Protocol (EAP) Chaining with Tunnel-based Extensible Authentication Protocol (TEAP) and Transport Layer Security (TLS), abbreviated EAP/TEAP-TLS: TEAP is a tunnel-based EAP method that establishes a secure tunnel and executes other EAP methods under the protection of that secured tunnel. ISE is used to validate user credentials and to send user sessions to the Cloud-Delivered Firewall Management Center. For more information, see How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS, on page 2292.

Note Before deploying policies related to a Microsoft Azure AD realm, see User limits for microsoft azure active directory realms, on page 2271.

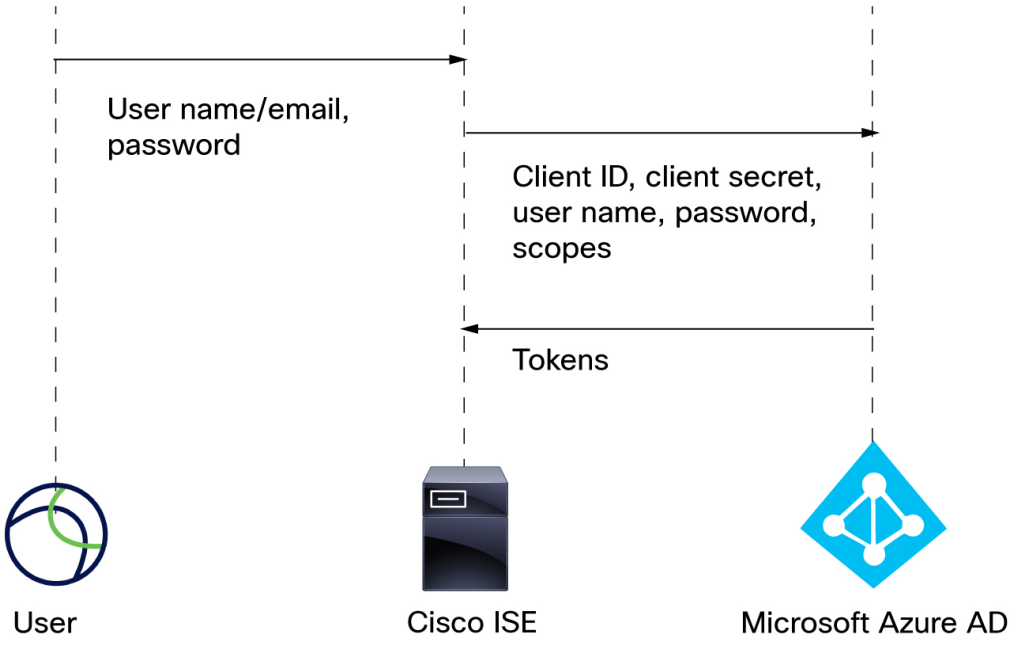
How Entra ID and Cisco ISE authentication with resource owned password credentials works

Summary

The key components involved in the Entra ID authentication with resource owned password credentials process are:

- User: Logs in with credentials using a VPN client
- Cisco Secure Client: VPN client that sends user credentials
- Entra ID (formerly Azure AD): Validates credentials and issues tokens
- Cisco ISE: Receives tokens and manages user sessions
- Cloud-Delivered Firewall Management Center: Receives user sessions from ISE

Workflow



These stages describe how Entra ID and Cisco ISE authentication with resource owned password credentials works:

1. The user logs in with a user name (or email address) and password using a VPN client like Cisco Secure Client.
2. The client ID, client secret, user name, password, and scopes are sent to Entra ID.
3. Tokens are sent from Entra ID to Cisco ISE, which sends user sessions to the Cloud-Delivered Firewall Management Center.

   For details about configuring Cisco ISE, see Configure ISE 3.0 REST ID with Azure Active Directory.

   Additional resource: Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials on learn.microsoft.com.

How Entra ID and Cisco ISE authentication works with TEAP/EAP-TLS

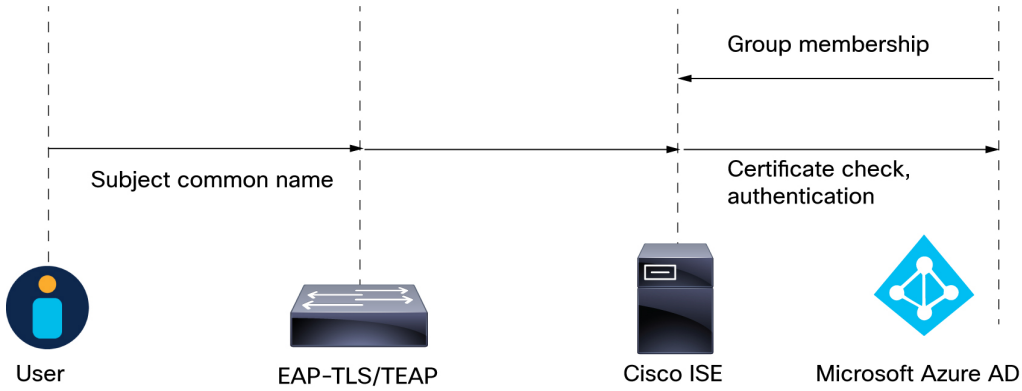
Tunnel Extensible Authentication Protocol (TEAP), defined by RFC7170, can be used with ISE and the Security Cloud Control for secure authentication.

Summary

The key components involved in this authentication process are:

- User device: Provides certificate through EAP-TLS or TEAP protocol
- ISE: Validates certificate and performs user lookup through Azure Graph API
- Entra ID: Provides user attributes and group information for authorization

Workflow



These stages describe how the authentication process works and are based on Configure Cisco ISE 3.2 EAP-TLS with Microsoft Azure Active Directory:

1. The user's certificate is sent to ISE through EAP-TLS or TEAP with EAP-TLS as the inner method.
2. ISE evaluates the user's certificate (validity period, trusted certificate authority, certificate revocation list, and so on).
3. ISE takes the certificate subject name (CN) and performs a look-up to the Azure Graph API to fetch the user's groups and other attributes. This is referred to by Azure as User Principal name (UPN).
4. ISE authorization policies are evaluated against the user's attributes returned from Azure.

Create a Microsoft Azure AD realm for passive authentication

Create and configure a Microsoft Azure AD (SAML) (now called *Entra ID*) realm for passive authentication use with the Security Cloud Control.

This configuration enables passive authentication between Microsoft Entra ID and your firewall system through proper realm setup and access control policy integration.

Procedure

Step 1 Configure Microsoft Entra ID.

Several configuration tasks are required, including setting up an event hub, giving your application permission to the Microsoft Graph API, and enabling the audit log.

See Configure Microsoft Entra ID for passive authentication, on page 2293.

Step 2 Configure Cisco ISE.

The way you configure ISE depends on how users authenticate with your system. For more information, see Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD, on page 2293.

Step 3 Create a Cisco ISE identity source.

The identity source enables ISE to communicate with the Cloud-Delivered Firewall Management Center.

Step 4 Get the information required to configure your Microsoft Azure AD realm.

This information includes client and tenant IDs, client secret, and other information store in Microsoft Entra ID.

Step 5 Configure and verify your realm.

Test the realm's configuration before you start to use it in access control policies.

Create a Microsoft Azure AD (SAML) realm as discussed in Create a Microsoft Azure AD (SAML) realm, on page 2289

Step 6 Create access control policies and rules using your Microsoft Azure AD (SAML) realm.

Unlike other types of realms, you do not need to create an identity policy or associate the identity policy with an access control policy.

See Creating a basic access control policy, on page 2004 and Create and edit access control rules, on page 2066.

What to do next

See How Entra ID and Cisco ISE authentication with resource owned password credentials works, on page 2291.

Configure Microsoft Entra ID for passive authentication

Configure Microsoft Entra ID to enable passive authentication and user group synchronization with the Security Cloud Control.

This task provides basic information about how to set up a Microsoft Entra ID (formerly called *Microsoft Azure Active Directory (AD)*) as a realm you can use with the Security Cloud Control. We expect you to already be familiar with Entra ID; if not, consult documentation or a support resource before you get started.

Procedure

Step 1 Grant your Entra ID application the Microsoft Graph permissions.

Grant your Entra ID application the following permissions to Microsoft Graph as discussed in Authorization and the Microsoft Graph Security API on the Microsoft site:

- Reader role
- User.Read.All permission
- Group.Read.All permission

This permission enables the Security Cloud Control to download users and groups from Entra ID the first time.

Required information from this step for setting up the Entra ID realm in the Security Cloud Control:

- Name of the app you registered
- **Application (client) ID**

- **Client secret**
- **Directory (tenant) ID**

Step 2 Set up an event hub.

Set up the event hub as discussed in Quickstart: Create an event hub using Azure portal on the Microsoft site. The Security Cloud Control uses the event hub audit log to download periodic updates to users and groups.

More information: Features and terminology in Azure Event Hubs.

Important
You must choose the **Standard** pricing tier or better. If you choose **Basic**, the realm cannot be used.

Required information from this step for setting up the Entra ID realm in the Security Cloud Control:

- Namespace **Name**
- **Connection string—primary key**
- Event Hub **Name**
- Consumer group **Name**

Step 3 Enable the audit log.

Enable the audit log as discussed in Tutorial: Stream Azure Active Directory logs to an Azure event hub on the Microsoft site.

Step 4 Configure Cisco ISE for Entra ID.

To send user session information to the Security Cloud Control, configure Cisco ISE for Entra ID as discussed in Configure ISE 3.0 REST ID with Azure Active Directory.

What to do next

See Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD, on page 2293.

Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD

In a Microsoft Azure AD (SAML)Microsoft Azure AD (now called *Entra ID*) realm, it's the responsibility of ISE to send user sessions (login, logout) to the Cloud-Delivered Firewall Management Center. This topic discusses how to set up ISE for use with an Azure AD realm.

Resource owned password credentials authentication

To use ISE with Microsoft Azure AD (SAML)Microsoft Azure ADMicrosoft Azure AD implemented through Representational State Transfer (REST) Identity (ID) service with the help of Resource Owner Password Credentials (ROPC), see Configure ISE 3.0 REST ID with Azure Active Directory.

TEAP/EAP-TLS

To use ISE with authorization policies based on Azure AD group membership and other user attributes with EAP-TLS or TEAP as the authentication protocols, see Configure Cisco ISE 3.2 EAP-TLS with Microsoft Azure Active Directory.

What to do next

Get required information For Your Microsoft Azure AD realm, on page 2296

Get required information For Your Microsoft Azure AD realm

This task enables you to collect the necessary credentials and configuration information from Microsoft Entra ID (formerly Azure AD) and Event Hubs to establish a SAML authentication realm in your Secure Firewall Management Center.

This task explains how to get the information required to set up a Microsoft Azure AD (now called *Entra ID*) realm in the Security Cloud Control. You might have already obtained this information when you set up Microsoft Entra ID as discussed in Configure Microsoft Entra ID for passive authentication, on page 2293.

Procedure

Step 1 Log in to https://portal.azure.com/ as a user with at least the Product Designer role.

Step 2 At the top of the page, click **Microsoft Entra ID**.

Step 3 In the left column, click **App Registrations**.

a) If necessary, filter the list of displayed apps to show the one you want to use.

b) Click the name of your app.



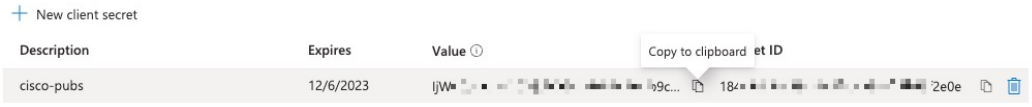
c) Click **Copy** ( ) next to the following values on this page and paste those values to a text file.

- **Application (Client) ID**
- **Directory (tenant) ID**

d) Click **Client Credentials**.

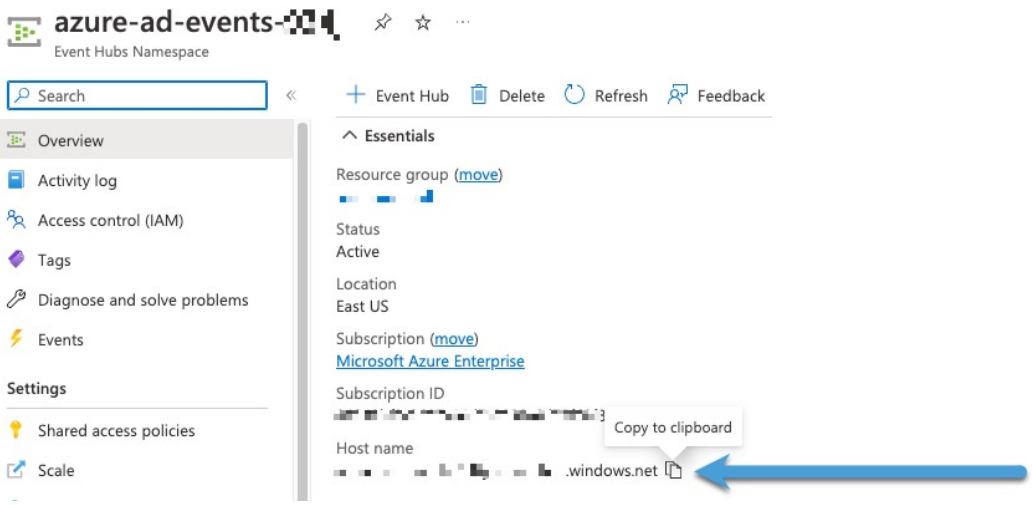
e) Unless you already know the client secret *value* (as opposed to the client *secret ID*), you must create a new client secret as follows:

1. Click **New Client Secret**.
2. Enter the required information in the provided fields.
3. Click **Add**.
4. Click **Copy** ( ) next to Value as the following figure shows.



Step 4 From https://portal.azure.com/, click **Event Hubs** > **(name of an event hub)**.

a) In the right pane, click **Copy** ( ) next to the value of **Host name** and paste the value to the clipboard. This is your event hub host name.

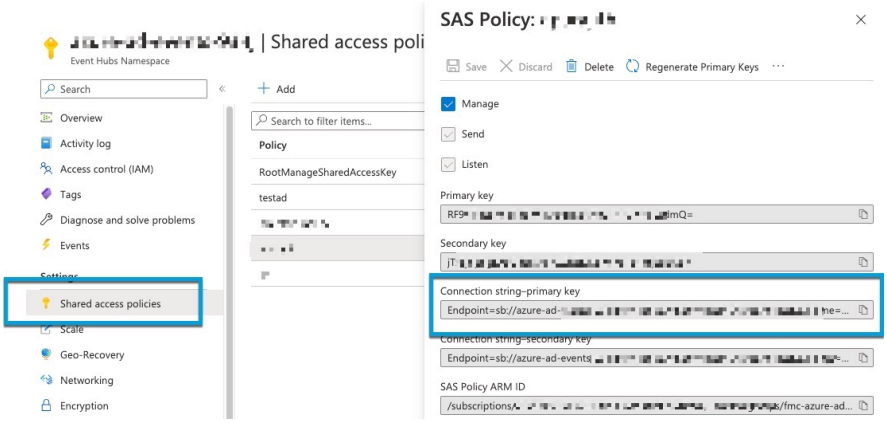


b) Write down or copy to a text file the name of the event hub (same as the **Event Hubs Namespace** at the top of the page).

c) In the left pane, under Settings, click **Shared access policies**.

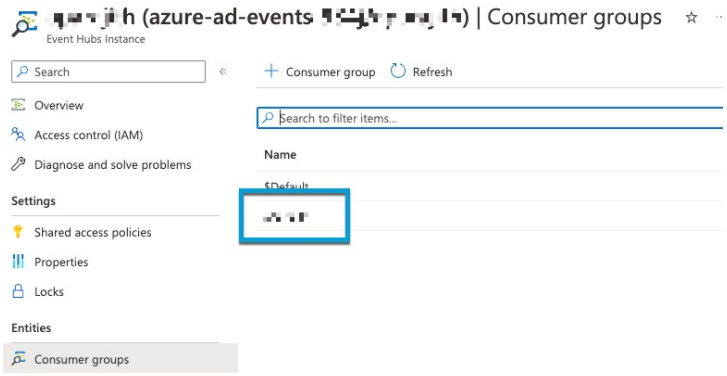
d) Click the name of a policy.

e) Click **Copy** ( ) next to **Connection string-primary key**.



Step 5 Click **Overview** > **Entities** > **Event Hubs** > **(name of an event hub)** > **Entities** > **Consumer Groups**.

Write down this value or copy it to the clipboard. This is your consumer group name.



Step 6 In the left pane, click **Overview**.

Step 7 Click **Copy** ( ) next to **Namespace**.



This is your event hubs topic name.

Create an Azure AD realm

This task creates a *realm* (a connection between the Cloud-Delivered Firewall Management Center and a Microsoft Azure AD realm) to enable user and identity control through access control policies.

Note To perform user and identity control with an Azure AD realm, you need only an access control policy with an associated Azure AD realm. You do not need to create an identity policy.

Before you begin

Complete all of the following tasks:

- Configure ISE as discussed in Configure Cisco ISE for Microsoft Azure AD (SAML)Microsoft Azure AD, on page 2293
- Create an ISE identity source as discussed in ISE/ISE-PIC configuration, on page 2382
- Get values required for the Azure AD realm as discussed in Get required information For Your Microsoft Azure AD realm, on page 2296.
- Configure Azure AD as discussed in Configure Microsoft Entra ID for passive authentication, on page 2293

If you enabled Change Management, you must approve all certificates used in this procedure. Open a new ticket or edit an existing one. For more information, see Create change management tickets, on page 175 and .Policies and objects that support change management, on page 169

Follow these steps to create an Azure AD realm:

Procedure

Step 1 Click **Administration** > **Integrations** > **Firewall Management Center** and choose **(select a management center)** > **Devices**

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 To create a new realm, click **Add Realm** > **Azure AD**.

Step 4 Enter the required information.

Item	Description
Name	
(Optional.) **Description**	
Client ID	Enter the information you found as discussed in Get required information For Your Microsoft Azure AD realm, on page 2296.
Client Secret	
Tenant ID	
Event Hubs Host Name	
Event Hub Name	
Event Hub Connection String	
(Optional.) **Excluded User Groups**	Enter one or more groups from which to *not* download users for identity control. Users in these groups will not be available for use in access control policies. Enter one group name per line followed by a line break. Group names are case-sensitive.

Step 5 Enter the values you found as discussed in Get required information For Your Microsoft Azure AD realm, on page 2296.

Step 6 Click **Test**.

Step 7 Fix any errors that are displayed in the test.

Step 8 Click **Save**.

What to do next

Create an access control policy and rule as discussed in Creating a basic access control policy, on page 2004.

Azure AD user session timeout

Set the Entra ID user session timeout for ISE or ISE-PIC to define how long users stay authenticated, ensuring network access control after timeout via the User Session Timeout page when editing an Entra ID realm.

The Azure AD user session timeout for ISE/ISE-PIC is available on the **User Session Timeout** page when you edit an Azure AD realm.

The default value 1440 minutes (24 hours). After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Cloud-Delivered Firewall Management Center as Unknown (except for **Failed Captive Portal Users**).

Create an LDAP realm or an Active Directory realm and realm directory

Before you begin

If you're using Kerberos authentication for captive portal, see the following section before you begin: Prerequisites for kerberos authentication, on page 2281 .

If you have enabled Change Management, you must open or edit, assign, and approve a ticket for each of the following objects before you can create the realm:

- If you're connecting securely to Microsoft AD or LDAP, the server's trusted certificate authority.
- The realm itself.

For more information, see Opening a ticket for configuration changes, on page 176 and Policies and objects that support change management, on page 169.

If you are managing devices with Security Cloud Control (Security Cloud Control), create a proxy sequence first as discussed in Create a proxy sequence, on page 2288 .



Important To reduce latency between the Cloud-Delivered Firewall Management Center and your Active Directory domain controller, we strongly recommend you configure a realm directory (that is, domain controller) that is as close as possible geographically to the Cloud-Delivered Firewall Management Center.

For example, if your Cloud-Delivered Firewall Management Center is in North America, configure a realm directory that is also in North America. Failure to do so can cause problems such as timeout downloading users and groups.

If you're setting up ISE/ISE-PIC without a realm, be aware there is a user session timeout that affects how users are seen by the Cloud-Delivered Firewall Management Center.

The following procedure enables you to create a *realm* (a connection between the Cloud-Delivered Firewall Management Center and an Active Directory realm) and a *directory* (a connection between the Cloud-Delivered Firewall Management Center and an LDAP server or an Active Directory domain controller).

(Recommended.) To connect securely from the Cloud-Delivered Firewall Management Center to your Active Directory server, first perform the following tasks:

- Export the active directory server's root certificate, on page 2329

- Find the active directory server's name, on page 2328

Microsoft has announced that Active Directory servers will start enforcing LDAP binding and LDAP signing in 2020. Microsoft is making these a requirement because when using default settings, an elevation of privilege vulnerability exists in Microsoft Windows that could allow a man-in-the-middle attacker to successfully forward an authentication request to a Windows LDAP server. For more information, refer to 2020 LDAP channel binding and LDAP signing requirement for Windows on the Microsoft support site.

For more information about realm and directory configuration fields, see Realm fields, on page 2281 and Realm directory and synchronize fields, on page 2285 .

A step-by-step example of setting up a realm with cross-domain trust is shown in Cross-domain trust setups with Cloud-Delivered Firewall Management Center, on page 2333 .

An Active Directory Global Catalog server is *not supported* as a realm directory. For more information about the Global Catalog Server, refer to Global Catalog on learn.microsoft.com.

Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different Microsoft AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

Procedure

Step 1 Log in to the Security Cloud Control. Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**

Step 2 To create a new realm, choose from **Add Realm** drop-down list.

Step 3 To perform other tasks (such as enable, disable, or delete a realm), see Manage a realm, on page 2341 .

Step 4 Enter realm information as discussed in Realm fields, on page 2281 .

Step 5 (Optional.) From the **Proxy** list, click a managed device or proxy sequence to communicate with ISE/ISE-PIC if Security Cloud Control is unable to do so.

Step 6 In the Directory Server Configuration section, enter directory information as discussed in Realm directory and synchronize fields, on page 2285 .

Step 7 (Optional.) To configure another domain for this realm, click **Add another directory** .

Step 8 Click **Configure Groups and Users** .

Information	Description
AD Primary Domain	Domain for the Active Directory server where users should be authenticated. For additional information, see Realm fields, on page 2281 . You must create the realm with the original domain name of the domain and not any alternative user principal name (UPN) suffixes of the domain. Otherwise, users and groups fail to download and identity policies will not be enforced. For example, if the original domain is `domain.example.com` and the alternative UPN name is`domain2.mydomain.com`, you must configure the realm to use `domain.example.com`. For more information about configuring an alternative UPN suffix, see a resource like Configuring Alternate Login ID on learn.microsoft.com.
Base DN	The directory tree on the server where the Cloud-Delivered Firewall Management Center should begin searching for user data.
Group DN	The directory tree on the server where the Cloud-Delivered Firewall Management Center should begin searching for group data.
Proxy	From the list, click one or more managed devices or a proxy sequence. These devices must be able to communicate with Active Directory or ISE/ISE-PIC to retrieve user data for identity policies.
Load Groups	Click to load groups from the Active Directory server. If no groups are displayed, enter or edit information in the **AD Primary Domain** , **Base DN** , and **Group DN** fields and click **Load Groups** . For more information about those fields, see Realm fields, on page 2281 .
Available Groups section	Limit the groups to use in policy by moving them to either the **Included Groups and Users** or **Excluded Groups and Users** list. Moving one group to the **Included Groups and Users** list, for example, allows that group only to be used in policy but excludes all other groups. Groups in the **Excluded Groups and Users** and the users they contain are excluded from user awareness and control. All other groups and users *are* available. For more information, see Realm directory and synchronize fields, on page 2285 .

Configure the Realm

Before you begin

You must have created a new realm and its groups and users, as described in the prior procedure.

Procedure

Step 1 Click the **Realm Configuration** tab.

Step 2 Enter **Group Attribute** , and (if you use Kerberos authentication for captive portal) enter **AD Join Username** and **AD Join Password** . For more information, see Realm directory and synchronize fields, on page 2285 .

Step 3 If you use Kerberos authentication, click **Test** . If the test fails, wait a short time and try again.

Step 4 Enter user session timeout values, in minutes, for **ISE/ISE-PIC Users** , **Terminal Server Agent Users** , **Captive Portal Users** , **Failed Captive Portal Users** , and **Guest Captive Portal Users** .

Step 5 When you're finished configuring the realm, click **Save** .

Connect securely to Active Directory or LDAP

To create a secure connection between an Active Directory or LDAP server and the Cloud-Delivered Firewall Management Center (which we strongly recommend), you must perform all of the following tasks:

- Export the server's *root* certificate.
- Import the root certificate into the Cloud-Delivered Firewall Management Center as a trusted CA certificate (**Objects** > **Object Management** > **PKI** > **Trusted CAs**).
- Find the server's fully qualified name.
- Create the realm directory.

Related Topics

Export the active directory server's root certificate, on page 2329
Find the active directory server's name, on page 2328
Create an LDAP realm or an Active Directory realm and realm directory, on page 2325

Find the active directory server's name

Find the fully qualified server name in Microsoft Active Directory that you need to configure a realm directory.

To configure a realm directory in the Cloud-Delivered Firewall Management Center, you must know the fully qualified server name, which you can find as discussed in the procedure that follows.

These tasks apply to Microsoft Active Directory only. If you use LDAP, consult an appropriate reference for the procedure.

Before you begin

You must log in to the Active Directory server as a user with sufficient privileges to view the computer's name.

Follow these steps to find the Active Directory server's name:

Procedure

Step 1 Log in to the Active Directory server.

Step 2 Click **Start**.

Step 3 Right-click **This PC**.

Step 4 Click **Properties**.

Step 5 Click **Advanced System Settings**.

Step 6 Click the **Computer Name** tab.

Step 7 Note the value of **Full computer name**.
You must enter this exact name when you configure the realm directory in the Cloud-Delivered Firewall Management Center.

What to do next

Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

Related Topics

Export the active directory server's root certificate, on page 2329

Export the active directory server's root certificate

Export the Active Directory server's root certificate to enable secure connection to the Cloud-Delivered Firewall Management Center for obtaining user identity information and authentication purposes.

The task that follows discusses how to export the Active Directory server's root certificate, which is required to connect securely to the Cloud-Delivered Firewall Management Center to obtain user identity information.

These tasks apply to Microsoft Active Directory only. If you use LDAP, consult an appropriate reference for the procedure.

Before you begin

You must know the name of your Active Directory server's root certificate. The root certificate might have the same name as the domain or the certificate might have a different name. The procedure that follows shows one way you can find the name; there could be other ways, however.

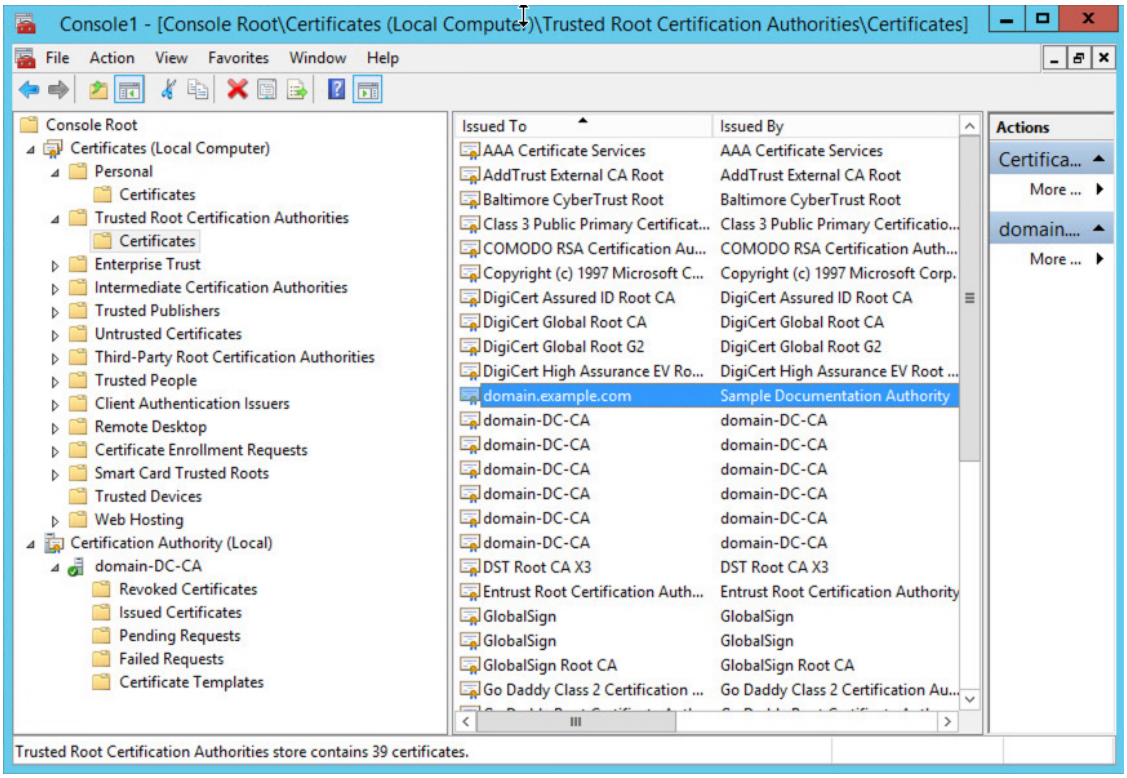
Follow these steps to export the Active Directory server's root certificate:

Procedure

Step 1 Find the name of the Active Directory Server's root certificate using the Microsoft Management Console.

This is one way to find the name of the Active Directory Server's root certificate; consult Microsoft documentation for more information:

a) Log in to the Active Directory server as a user with privileges to run the Microsoft Management Console.
b) Click **Start** and enter **mmc**.
c) Click **File** > **Add/Remove Snap-in**
d) From the Available Snap-ins list in the left pane, click **Certificates (local)**.
e) Click **Add**.
f) At the Certificates snap-in dialog box, click **Computer Account** and click **Next**.
g) At the Select Computer dialog box, click **Local Computer** and click **Finish**.
h) *Windows Server 2012 only*. Repeat the preceding steps to add the Certification Authority snap-in.
i) Click **Console Root** > **Trusted Certification Authorities** > **Certificates**.
The server's trusted certificates are displayed in the right pane. The following figure is only an example for Windows Server 2012; yours will probably look different.



Step 2 Export the certificate using the **certutil** command.

This is only one way to export the certificate. It is a convenient way to export the certificate, especially if you can run a web browser and connect to the Cloud-Delivered Firewall Management Center from the Active Directory server.

a) Click **Start** and enter **cmd**.
b) Enter the command **certutil -CA.cert** *certificate-name*.
The server's certificate is displayed on the screen.
c) Copy the entire certificate to the clipboard, starting with **-----BEGIN CERTIFICATE-----** and ending with **-----END CERTIFICATE-----** (including those strings).

What to do next

Import the Active Directory server's certificate into the Cloud-Delivered Firewall Management Center as a trusted CA certificate as discussed in Add a trusted CA object, on page 2993.

Related Topics

Find the active directory server's name, on page 2328

Synchronize users and groups

This task synchronizes users and groups from Active Directory domains to enable their use in identity policies.

Synchronizing users and groups means the Cloud-Delivered Firewall Management Center queries the realms and directories you configured for groups and users in those groups. All users the Cloud-Delivered Firewall Management Center finds can be used in identity policies.

If issues are found, you most likely need to add a realm that contains users and groups the Cloud-Delivered Firewall Management Center cannot load. For details, see Realms and trusted domains, on page 2275.

Before you begin

Create a Cloud-Delivered Firewall Management Center *realm* for each Active Directory domain and a Cloud-Delivered Firewall Management Center *directory* for each Active Director domain controller in each forest. See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

You must create a realm only for domains that have users you want to use in user control.

You can nest Microsoft AD groups and the Cloud-Delivered Firewall Management Center downloads those groups and the users they contain. You can optionally restrict which groups and users get downloaded as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

You must create the realm with the original domain name of the domain and not any alternative user principal name (UPN) suffixes of the domain. Otherwise, users and groups fail to download and identity policies will not be enforced. For example, if the original domain is `domain.example.com` and the alternative UPN name is`domain2.mydomain.com`, you must configure the realm to use `domain.example.com`. For more information about configuring an alternative UPN suffix, see a resource like Configuring Alternate Login ID on learn.microsoft.com.

Follow these steps to synchronize users and groups:

Procedure

Step 1 Log in to the Security Cloud Control.

Step 2 Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 3 Next to each realm, click **Download** ( ).

Step 4 To see the results, click the **Sync Results** tab.
The Realms column indicates whether or not there were issues synchronizing users and groups in Active Directory forests. Look for the following indicators next to each realm.

Indicator in Realms column	Meaning
(nothing)	All users and groups synchronized without error. No action is necessary.
Yellow Triangle ( )	There were issues synchronizing users and groups. Make sure you added a realm for each Active Directory domain and a directory for each Active Directory domain controller. For more details, see Troubleshoot cross-domain trust, on page 2345.

Create a realm sequence

Creating a realm sequence provides an ordered list of realms the system searches when it applies identity policy, allowing you to specify the priority order for user authentication across multiple Active Directory servers.

The following procedure enables you to create a realm sequence, which is an ordered list of realms the system searches when it applies identity policy. You add a realm sequence to an identity rule exactly the same way as you add a realm; the difference is that the system searches all the realms in the order specified in the realm sequence when applying an identity policy.

Before you begin

You must create and enable at least two realms, each corresponding to a connection with an Active Directory server. You cannot create realm sequences for LDAP realms.

Create a realm as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

Follow these steps to create a realm sequence:

Procedure

Step 1 Log in to the Security Cloud Control and click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 2 Under the **Realm Sequences** tab, click **Add Sequence**.

Step 3 In the **Name** field, enter a name to identify the realm sequence.

Step 4 (Optional.) In the **Description** field, enter a description for the realm sequence.

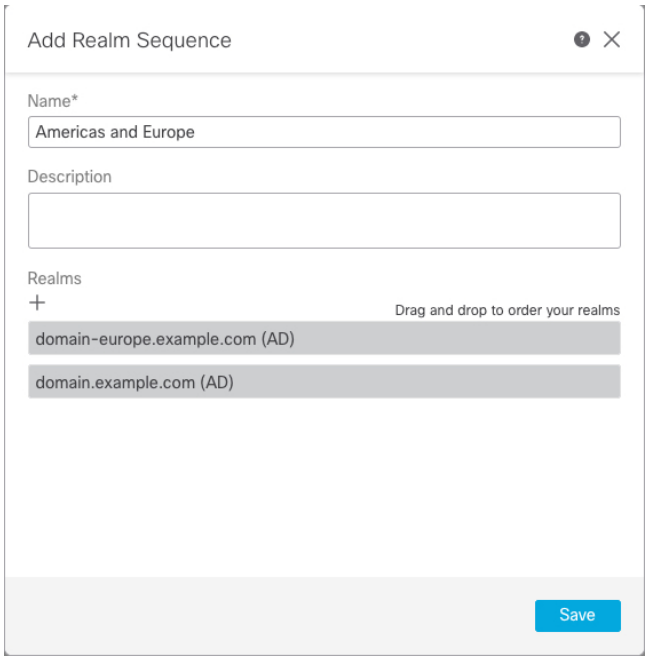
Step 5 Under Realms, click **Add** (+).

Step 6 Click the name of each realm to add to the sequence.

To narrow your search, enter all or part of a realm name in **Filter** field.

Click **OK**.

Step 7 In the Add Realm Sequence dialog box, drag and drop the realms in the order in which you want the system to search for them.
The following figure shows an example of a realm sequence consisting of two realms. The **domain-europe.example.com** realm will be searched for users before the **domain.example.com** realm.



Step 8 Click **Save**.

What to do next

See Create an identity policy, on page 2507.

Cross-domain trust setups with Cloud-Delivered Firewall Management Center

A cross-domain trust setup allows discovery and use of users and groups from multiple domains in access control policies.

This is an introduction to several topics that walk you through configuring the Cloud-Delivered Firewall Management Center with two realms having cross-domain trust.

Example setup configuration

This step-by-step example involves two forests: **forest.example.com** and **eastforest.example.com**. The forests are configured so that certain users and groups in each forest can be authenticated by Microsoft AD in the other forest.

The example setup uses this configuration:



Using this example, you would set up the Cloud-Delivered Firewall Management Center with these components:

- Realm and directory for any domain in **forest.example.com** that contains users you want to control with access control policy
- Realm and directory for any domain in **eastforest.example.com** that contains users you want to control with access control policy

Each realm in the example has one domain controller, which is configured in the Cloud-Delivered Firewall Management Center as a directory. The directories in this example are configured as follows:

- **forest.example.com**
  - Base distinguished name (DN) for users: **ou=UsersWest,dc=forest,dc=example,dc=com**
  - Base DN for groups: **ou=EngineringWest,dc=forest,dc=example,dc=com**
- **eastforest.example.com**
  - Base DN for users: **ou=EastUsers,dc=eastforest,dc=example,dc=com**
  - Base DN for groups: **ou=EastEngineering,dc=eastforest,dc=example,dc=com**

Related Topics

Configure Security Cloud Control for cross-domain-trust step 1: Configure realms and directories, on page 2334

Configure Security Cloud Control for cross-domain-trust step 1: Configure realms and directories

This task configures the Cloud-Delivered Firewall Management Center to recognize Active Directory servers in a cross-domain trust relationship, enabling the discovery and use of users and groups from multiple domains in access control policies.

This is the first task in a step-by-step procedure that explains how to configure the Cloud-Delivered Firewall Management Center to recognize Active Directory servers configured in a cross-domain trust relationship, which is an increasingly common configuration for enterprise organizations. For an overview of this sample configuration, see Cross-domain trust setups with Cloud-Delivered Firewall Management Center, on page 2333.

If you set up the system with one realm for each domain and one directory for each domain controller, the system can discover up to 100,000 foreign security principals (users and groups). If these foreign security principals match a user downloaded in another realm, then they can be used in access control policy.

Before you begin

You must configure Microsoft Active Directory servers in a cross-domain trust relationship; see Realms and trusted domains, on page 2275 for more information.

If you authenticate users with LDAP, you *cannot* use this procedure.

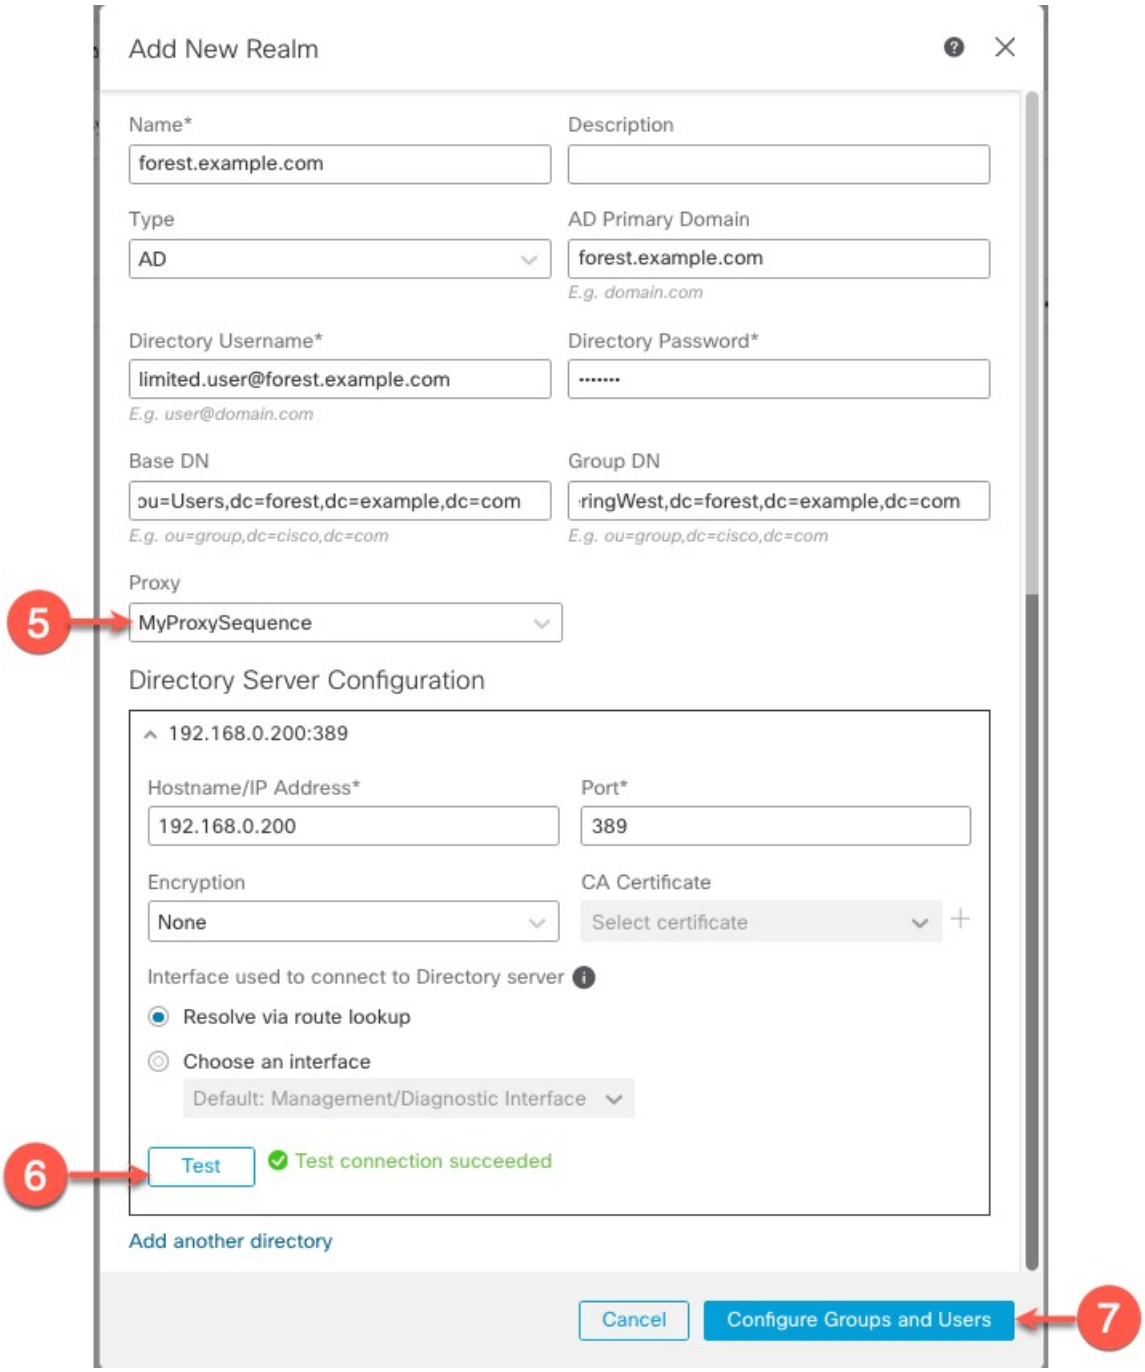
Procedure

Step 1 Log in to the Security Cloud Control.

Step 2 Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 3 Click **Add Realm** > **Active Directory/LDAP**.

Step 4 Enter the following information to configure **forest.example.com**.



Note
The **Directory Username** can be any user in the Active Directory domain; no special permissions are required.

The **Interface used to connect to Directory server** can be any interface that can connect to the Active Directory server.

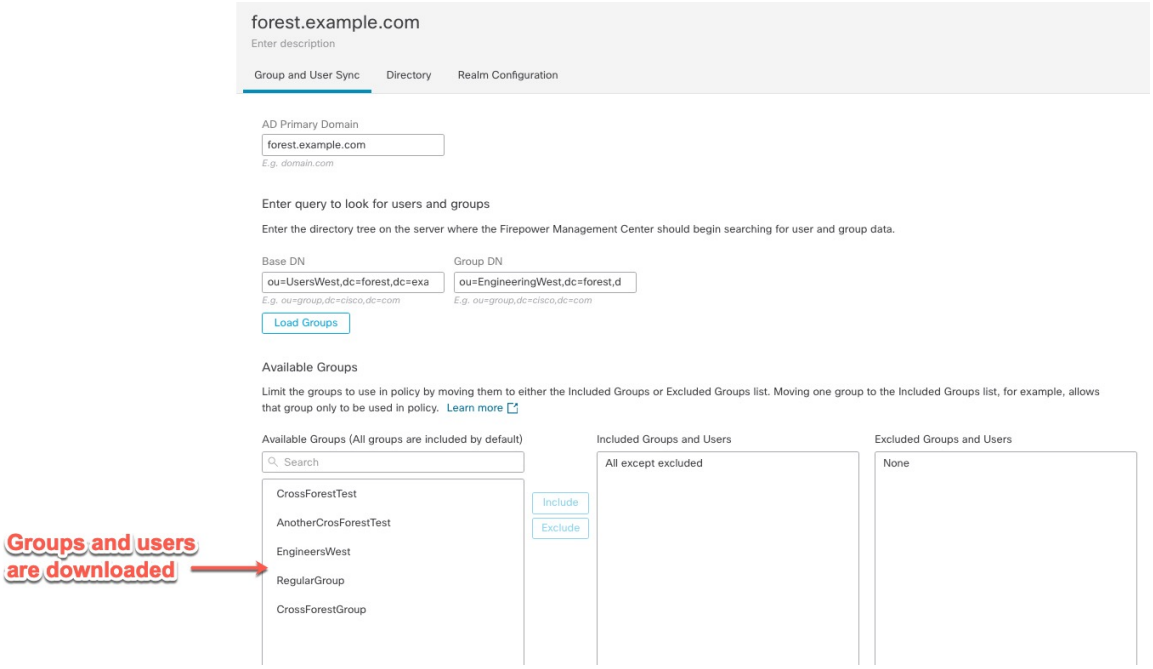
Step 5 **Proxy** is an optional managed device or proxy sequence to communicate with ISE/ISE-PIC if Security Cloud Control is unable to do so.

For example, your Security Cloud Control might be in a public cloud but the ISE/ISE-PIC server might be on an internal intranet.

Step 6 Click **Test** and make sure the test succeeds before you continue.

Step 7 Click **Configure Groups and Users**.

Step 8 Verify your configuration was successful by checking that the next page is displayed similar to the following.



Note
If groups and users were not downloaded, verify the values in the **Base DN** and **Groups DN** fields and click **Load Groups**.

There are other optional configurations available on this page; for more information about them, see Realm fields, on page 2281 and Realm directory and synchronize fields, on page 2285.

Step 9 If you made changes on this page or tab pages, click **Save**.

Create another realm

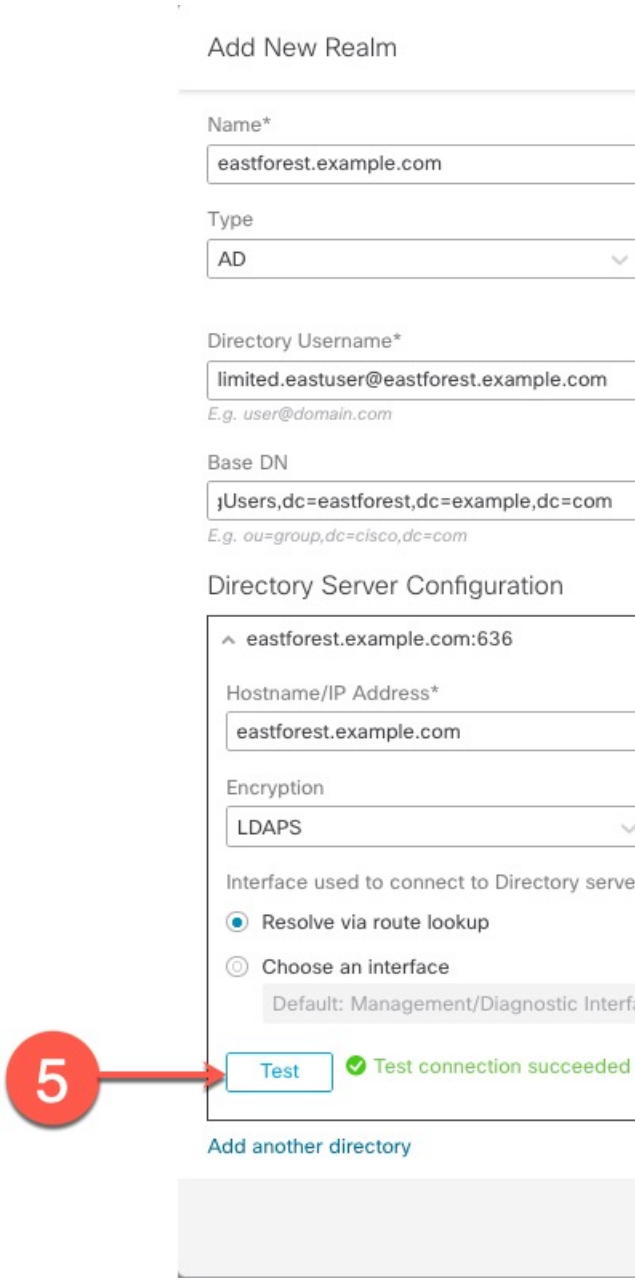
Before you begin

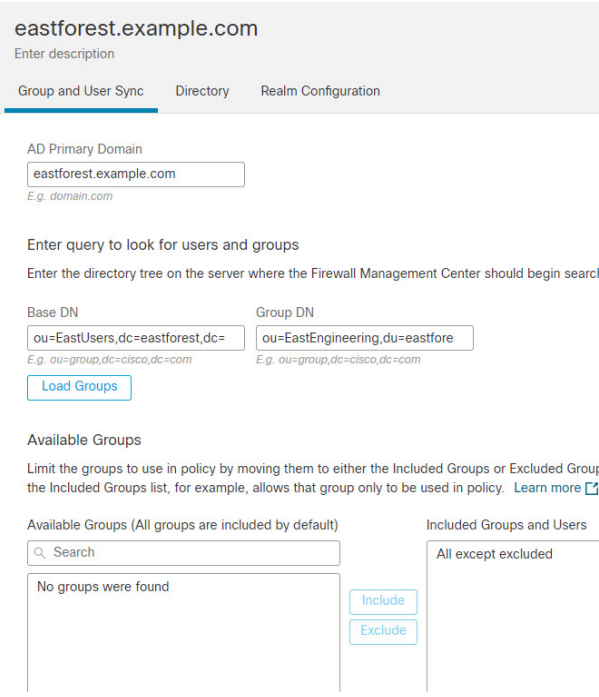
Complete the configuration of your first realm as in the prior task.

Follow these steps to create another realm.

Procedure

	Command or Action	Purpose
Step 1	Click **Integration** > **Other Integrations** > **Realms** > **Realms**.	

	Command or Action	Purpose
Step 2	Click **Add Realm**.	
Step 3	Enter the following information to configure **eastforest.example.com**.	
Step 4	Click **Test** and make sure the test succeeds before you continue.	
Step 5	Click **Configure Groups and Users**.	

	Command or Action	Purpose
Step 6	Verify your configuration was successful by checking that the next page is displayed similar to the following.	

You have successfully configured two realms (**forest.example.com** and **eastforest.example.com**) with their respective directories.

Configure the Security Cloud Control for cross-domain trust Step 2: Synchronize users and groups

This task allows you to download users and groups from Active Directory servers with cross-domain trust relationships to synchronize identity data in the Secure Firewall Management Center and resolve download issues by verifying directory setup and device communication.

After you configure two or more Active Directory servers that have a cross-domain trust relationship, you must download users and groups. That process exposes possible issues with the Active Directory configuration (for example, groups or users downloaded for one Active Directory domain but not the other).

Before you begin

Make sure you have performed the tasks discussed in Configure Security Cloud Control for cross-domain-trust step 1: Configure realms and directories, on page 2334.

Procedure

Step 1 Log in to the Security Cloud Control.

Step 2 Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 3 At the end of the row of any realm in the cross-domain trust, click **Download Now**( ), then click **Yes**.

Step 4 Click **Check Mark** ( ) (Notifications) > **Tasks**.

If groups and users fail to download, try again. If subsequent attempts fail, review your realm and directory setup as discussed in Realm fields, on page 2281 and Realm directory and synchronize fields, on page 2285.

If you're using a proxy or proxy sequence, make sure all managed devices can communicate with Active Directory or ISE/ISE-PIC. If more than one managed device can communicate with ISE/ISE-PIC, we strongly recommend you set up a proxy sequence for the realm as discussed in Create a proxy sequence, on page 2288

Step 5 Click **Integration** > **Other Integrations** > **Realms** > **Sync Results**.

Related Topics

Configure the Security Cloud Control for cross-domain trust step 3: Resolve issues, on page 2340

Configure the Security Cloud Control for cross-domain trust step 3: Resolve issues

This task enables you to make sure users and groups are downloaded without errors during cross-domain trust setup.

The final step in setting up cross-domain trust in the Cloud-Delivered Firewall Management Center is to make sure users and groups are downloaded without errors. A typical reason why users and groups do not download properly is that the realms to which they belong have not been downloaded to the Cloud-Delivered Firewall Management Center.

This topic discusses how to diagnose that a group referred in one forest to cannot be downloaded because the realm is not configured to find the group in the domain controller hierarchy.

Procedure

Step 1 Log in to the Security Cloud Control.

Step 2 Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

In the Realms column, if **Yellow Triangle** ( ) is displayed next to the name of a realm, you have issues that must be resolved. If not, your results are configured properly and you can quit.

Step 3 Download users and groups again from the realms that display issues.

a) Click the **Realms** tab.

b) Click (Download Now), then click **Yes**.

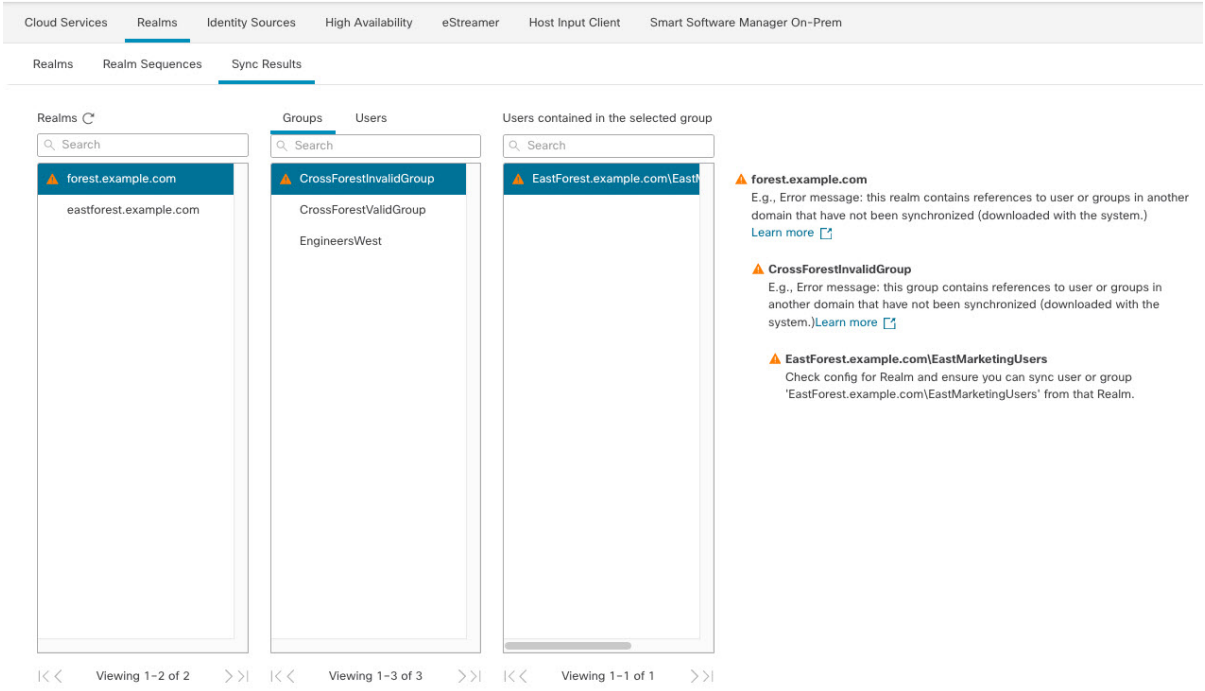
Step 4 Click the **Sync Results** tab page.

If the **Yellow Triangle** ( ) is displayed in the Realms column, click **Yellow Triangle** ( ) next to the realm that has issues.

Step 5 In the middle column, click either **Groups** or **Users** to find more information.

Step 6 In the Groups or Users tab page, click **Yellow Triangle** ( ) to display more information.

The right column should display enough information you can isolate the source of the issue.



In the preceding example, **forest.example.com** includes a cross-domain group **CrossForestInvalidGroup** that contains another group **EastMarketingUsers** that was not downloaded by the Cloud-Delivered Firewall Management Center. If, after synchronizing the **eastforest.example.com** realm again, the error does not resolve, it likely means that the Active Directory domain controller does not include **EastMarketingUsers**.

To resolve this issue, you can:

- Remove the **EastMarketingUsers** from **CrossForestInvalidGroup**, synchronize the **forest.example.com** realm again, and recheck.
- Remove the **ou=EastEngineering** value from the **Group DN** of the **eastforest.example.com** realm, which causes the Cloud-Delivered Firewall Management Center to retrieve groups from the highest level in the Active Directory hierarchy, synchronize **eastforest.example.com**, and recheck.

Manage a realm

This task enables you to perform various maintenance operations on realms to manage authentication and user access control.

This section discusses how to perform various maintenance tasks for a realm using controls on the Realms page:

Procedure

Step 1 Log in to the Security Cloud Control. Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 2 To delete a realm, click **Delete** ( ).

Step 3 To edit a realm, click **Edit** ( ) next to the realm and make changes as described in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

Step 4 To enable a realm, slide **State** to the right; to disable a realm, slide it to the left.

Step 5 To download users and user groups, click **Download** ( ).

Step 6 To copy a realm, click **Copy** ( ).

Step 7 To compare realms, see Compare realms, on page 2342.

Compare realms

Use this task to compare identity realms in the Secure Firewall Management Center to review differences between selected realms, generate comparison reports, and analyze changes to support identity integration management.

You must be an Admin, Access Admin, Network Admin, or Security Approver to perform this task.

Procedure

Step 1 Log in to the Security Cloud Control. Click **Policies** > **Firewall Threat Defense** > **Integrations** > **Other Integrations** > **Realms**.

Step 2 Click **Compare Realms**.

Step 3 Choose **Compare Realm** from the **Compare Against** list.

Step 4 Choose the realms you want to compare from the **Realm A** and **Realm B** lists.

Step 5 Click **OK**.

Step 6 To navigate individually through changes, click **Previous** or **Next** above the title bar.

Step 7 (Optional.) Click **Comparison Report** to generate the realm comparison report.

Step 8 (Optional.) Click **New Comparison** to generate a new realm comparison view.

Troubleshoot realms and user downloads

If you notice unexpected server connection behavior, consider tuning your realm configuration, device settings, or server settings. For other related troubleshooting information, see:

- Troubleshoot Cisco ISE/ISE-PIC or Cisco TrustSec issues, on page 2399
- Troubleshoot the TS agent identity source, on page 2502
- Troubleshoot the captive portal identity source, on page 2369
- Troubleshooting guidelines for remote access VPN identity source, on page 2497
- Troubleshoot user control, on page 2526

Users are not downloaded

Possible causes include:

- If you have the realm **Type** configured incorrectly, users and groups cannot be downloaded because of a mismatch between the attribute the system expects and what the repository provides. For example, if you configure **Type** as **LDAP** for a Microsoft Active Directory realm, the system expects the `uid` attribute, which is set to `none` on Active Directory. (Active Directory repositories use `sAMAccountName` for the user ID.)

  Solution: Set the realm **Type** field appropriately: **AD** for Microsoft Active Directory or **LDAP** for another supported LDAP repository.

- Users in Active Directory groups that have special characters in the group or organizational unit name might not be available for identity policy rules. For example, if a group or organizational unit name contains the characters asterisk (*), equals (=), or backslash (\), users in those groups are not downloaded and can't be used for identity policies.

  Solution: Remove special characters from the group or organizational unit name.

- If your managed device connects to Cloud-Delivered Firewall Management Center through its data interface, and a LDAP or Active Directory (AD) server is behind this managed device, the managed device acts as a proxy for the AD server traffic. Because the AD server is in the internal network, it might drop some request packets. This can cause realm download or test connection to fail.

  Solution: Create a flex config object with **ip-clien** *interface-name* and deploy it to to the device. Then download the users and groups.



Important To reduce latency between the Cloud-Delivered Firewall Management Center and your Active Directory domain controller, we strongly recommend you configure a realm directory (that is, domain controller) that is as close as possible geographically to the Cloud-Delivered Firewall Management Center.

For example, if your Cloud-Delivered Firewall Management Center is in North America, configure a realm directory that is also in North America. Failure to do so can cause problems such as timeout downloading users and groups.

Not all users in a realm are downloaded

Possible causes include:

- If you attempt to download more than the maximum number of users in any one realm, the download stops at the maximum number of users and a health alert is displayed. User download limits are set per Cloud-Delivered Firewall Management Center model.

- Every user must be a member of a group. Users that are members of no groups do not get downloaded.

Access control policy doesn't match group membership

This solution applies to an AD domain that is in a trust relationship with other AD domains. In this discussion, *external domain* means a domain other than the one to which the user logs in.

If a user belongs to a group defined in a trusted external domain, the Cloud-Delivered Firewall Management Center doesn't track membership in the external domain. For example, consider this scenario:

- Domain controllers 1 and 2 trust each other
- Group A is defined on domain controller 2
- User `mparvinder` in controller 1 is a member of Group A

Even though user `mparvinder` is in Group A, the Cloud-Delivered Firewall Management Center access control policy rules specifying membership Group A don't match.

Solution: Create a similar group in domain controller 1 that contains has all domain 1 accounts that belong to group A. Change the access control policy rule to match any member of Group A or Group B.

Access control policy doesn't match child domain membership

If a user belongs to a domain that is child of parent domain, Firepower doesn't track the parent/child relationships between domains. For example, consider this scenario:

- Domain `child.parent.com` is child of domain `parent.com`
- User `mparvinder` is defined in `child.parent.com`

Even though user `mparvinder` is in a child domain, the Firepower access control policy matching the `parent.com` don't match `mparvinder` in the `child.parent.com` domain.

Solution: Change the access control policy rule to match membership in either `parent.com` or `child.parent.com`.

Realm or realm directory test fails

The **Test** button on the directory page sends an LDAP query to the hostname or IP address you entered. If it fails, check these items:

- The **Hostname** you entered resolves to the IP address of an LDAP server or Active Directory domain controller.
- The **IP Address** you entered is valid.

The **Test AD Join** button on the realm configuration page verifies these items:

- DNS resolves the **AD Primary Domain** to an LDAP server or Active Directory domain controller's IP address.
- The **AD Join Username** and **AD Join Password** are correct.

  AD Join Username must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).

- The user has sufficient privileges to create a computer in the domain and join the Cloud-Delivered Firewall Management Center to the domain as a Domain Computer.

User timeouts are occurring at unexpected times

If you notice the system performing user timeouts at unexpected intervals, confirm that the time on your ISE/ISE-PIC server is synchronized with the time on the Cloud-Delivered Firewall Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.

If you notice the system performing user timeouts at unexpected intervals, confirm that the time on your ISE/ISE-PIC, or TS Agent server is synchronized with the time on the Cloud-Delivered Firewall Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.

User data for previously-unseen ISE/ISE-PIC users is not displaying in the web interface

After the system detects activity from an ISE/ISE-PIC or TS Agent user whose data is not yet in the database, the system retrieves information about them from the server. In some cases, the system requires additional time to successfully retrieve this information from Microsoft Windows servers. Until the data retrieval succeeds, activity seen by the ISE/ISE-PIC, or TS Agent user is *not* displayed in the web interface.

Note that this can also prevent the system from handling the user's traffic using access control rules.

User data in events is unexpected

If you notice user or user activity events contain unexpected IP addresses, check your realms. The system does not support configuring multiple realms with the same **AD Primary Domain** value.

Users originating from terminal server logins are not uniquely identified by the system

If your deployment includes a terminal server and you have a realm configured for one or more servers connected to the terminal server, you must deploy the Cisco Terminal Services (TS) Agent to accurately report user logins in terminal server environments. When installed and configured, the TS Agent assigns unique ports to individual users so the system can uniquely identify those users in the web interface.

For more information about the TS Agent, see the *Cisco Terminal Services (TS) Agent Guide*.

Troubleshoot cross-domain trust

Troubleshoot user and group synchronization issues in the Cloud-Delivered Firewall Management Center configuration for cross-domain trust. This reference provides information on typical issues, error interpretation, and resolution steps for Active Directory forest synchronization problems.

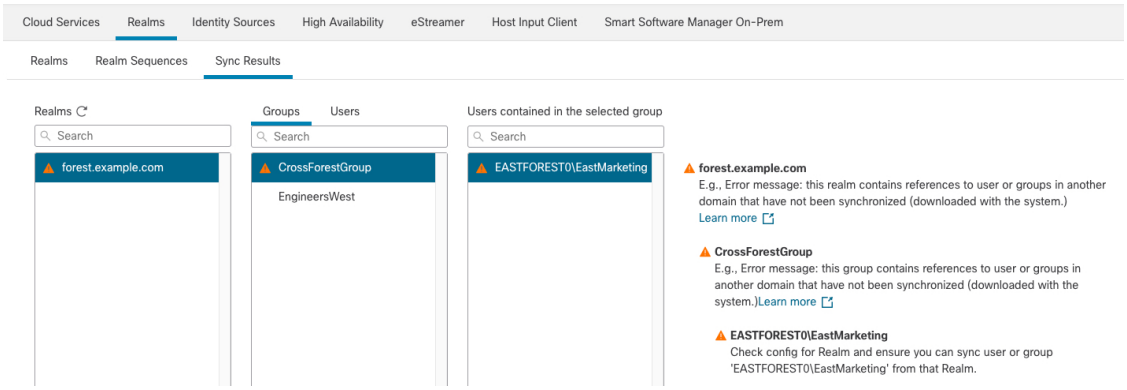
Typical issues

Typical issues with troubleshooting the Cloud-Delivered Firewall Management Center configuration for cross-domain trust include these items:

- Not adding realms or directories for all forests that have shared groups.
- Configure a realm to exclude users from being downloaded and those users are referenced in a group in a different realm.
- Certain temporary issues.

Understand the issues

If there are issues with the Firewall Management Center being able to synchronize users and groups with your Active Directory forests, the Sync Results tab page is displayed similar to this.



This table explains how to interpret the information.

Column	Meaning
Realms	Displays all realms configured in the system. Click **Refresh** ( ) to update the list of realms. **Yellow Triangle** ( ) is displayed to indicate issues in the realm. Nothing is displayed next to a realm if all users and groups synchronized successfully.
Groups	Click **Groups** to display all groups in the realm. As with realms, **Yellow Triangle** ( ) is displayed to indicate issues. Click **Yellow Triangle** ( ) to see more detail about the issue.
Users	Click **Users** to display all users, sorted by group.
Users contained in the selected group	Displays all users in the group you selected in the Groups column. Clicking **Yellow Triangle** ( ) displays more information to the right of the table.
Groups that contain selected user	Displays all groups the selected user belongs to. Clicking **Yellow Triangle** ( ) displays more information to the right of the table.

Column	Meaning
Error detail information (displayed to the right of the table).	The system displays the NetBIOS forest name and group name it could not synchronize. Typical reasons the system cannot synchronize these users and groups follow: • **Problem**: The forest containing the groups and users do not have corresponding realms configured in the Firewall Management Center. **Solution**: Add a realm for the forest that contains the group as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325. • **Problem**: You excluded groups from being downloaded to the Firewall Management Center. **Solution**: Click the **Realms** tab page, click **Edit** (), then move the indicated group or user from the **Excluded Groups and Users** list.

Try downloading users and groups again

If there is a possibility the issues are temporary, download users and groups for all realms.

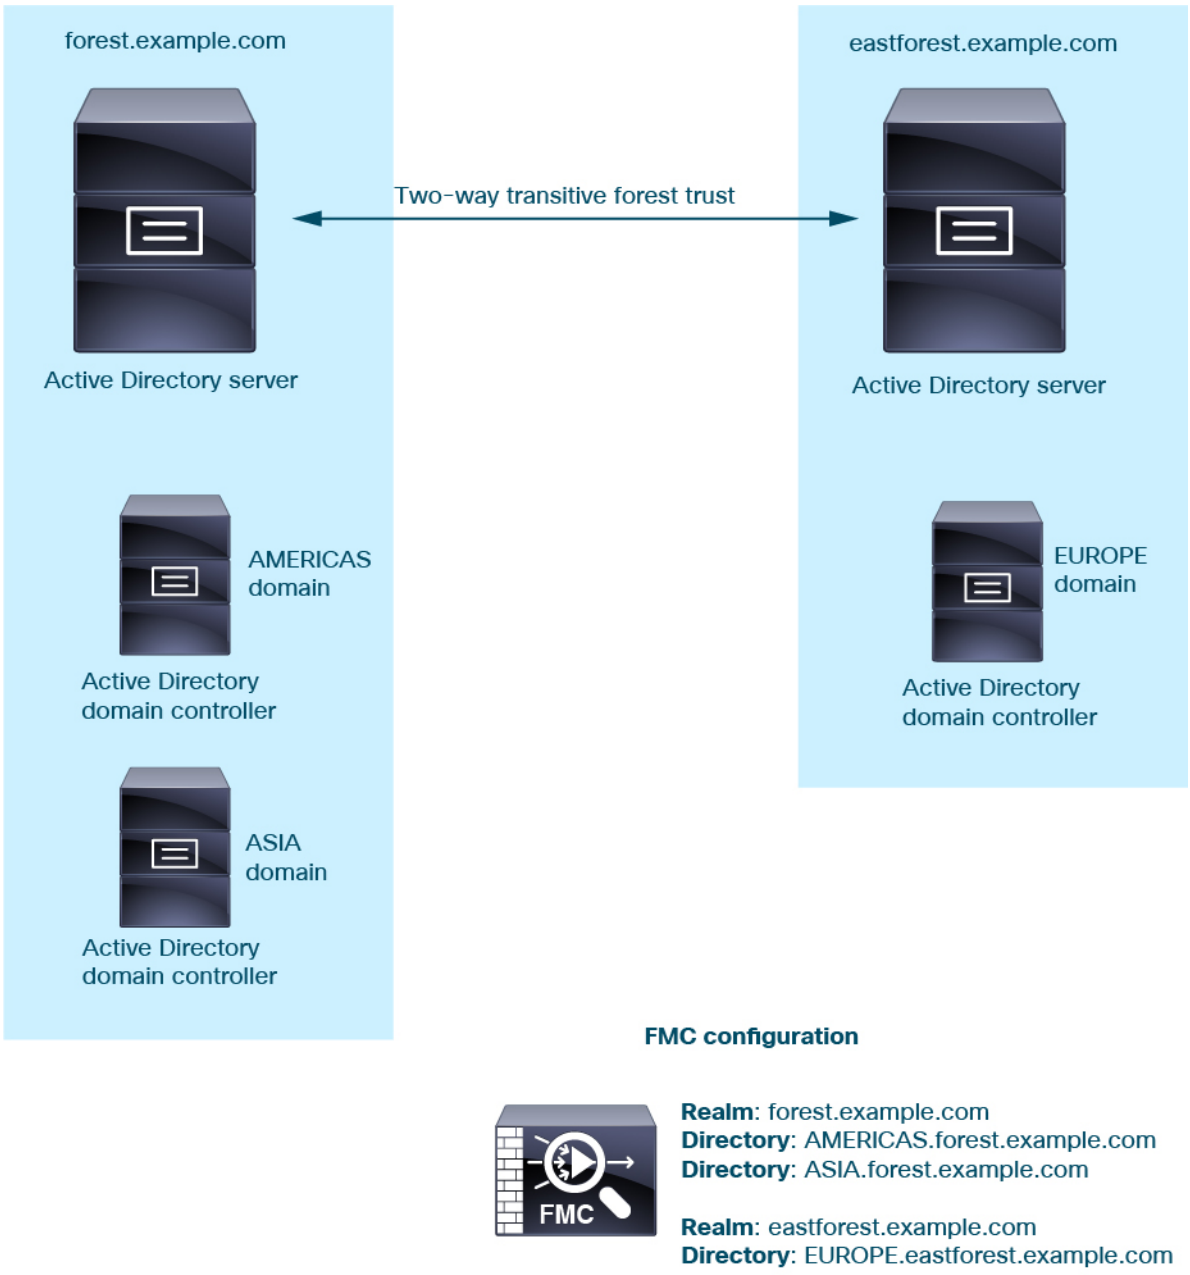
1. If you haven't done so already, log in to the Firewall Management Center.
2. Click **Integration** > **Other Integrations** > **Realms** > **Realms**.
3. Click **Download** ().
4. Click the **Sync Results** tab page.
5. If no indicator is displayed for entries in the Realms column, the issues have been resolved.

Add a realm for all forests

Make sure you configured:

- Firewall Management Center realm for each forest that has users you want to use in identity policies.
- Firewall Management Center directory for each domain controller in that forest with users you want to use in identity polices.

This figure shows an example.

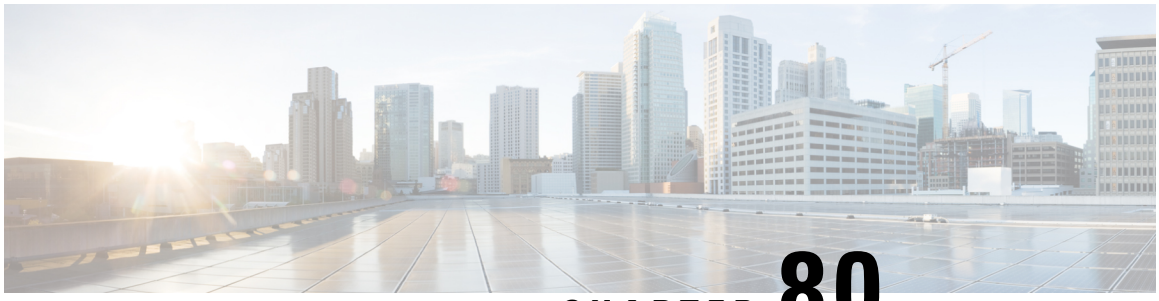


History for realms

This reference provides the version history and implementation details for realm-related features across different versions of the Cisco security platform.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Microsoft Azure Active Directory (SAML) realms.	November 8, 2024	7.4.0	You can now use Microsoft Azure Active Directory (AD) realms for active and passive authentication: • Active authentication using Azure AD: Use Azure AD as a captive portal. • Passive authentication using Cisco ISE (introduced in Version 7.4.0): The Cloud-Delivered Firewall Management Center gets groups from Azure AD and logged-in user session data from ISE. We use SAML (Security Assertion Markup Language) to establish a trust relationship between a service provider (the devices that handle authentication requests) and an identity provider (Azure AD). For upgraded Cloud-Delivered Firewall Management Centers, existing Azure AD realms are displayed as SAML - Azure AD realms. **Upgrade impact**. If you had a Microsoft Azure AD realm configured before the upgrade, it is displayed as a SAML - Azure AD realm configured for passive authentication. All previous user session data is preserved. New/modified screens: **Integration** > **Other Integrations** > **Realms** > **Add Realm** > **SAML - Azure AD** New/modified CLI commands: none See: Create a Microsoft Azure AD (SAML) Realm.
Microsoft Azure Active Directory (AD) realms.	January 18, 2023	7.4.0	You can use a Microsoft Azure Active Directory (AD) realm with ISE to authenticate users and get user sessions for user control. New/modified screens: **System** > **Integration** > **Realms** > **Add Realm** > **Azure AD**
Proxy sequences.	Any	7.2.0	Similar to a realm sequence, a proxy sequence is one or more managed devices that can communicate with Security Cloud Control in the event Security Cloud Control cannot communicate with the LDAP or Active Directory server. New/modified screens: **Integration** > **Other Integrations** > **Realms** > **Proxy Sequence**
Cross-domain trust for Active Directory domains.	Any	7.0.0	A grouping of Microsoft Active Directory (AD) domains that trust each other is commonly referred to as a *forest*. This trust relationship can enable domains to access each other's resources in different ways. For example, a user account defined in domain A can be marked as a member of a group defined in domain B. The Cloud-Delivered Firewall Management Center can get users from Active Directory forests for identity rules.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Realm sequences.	Any	6.7.0	A *realm sequence* is an ordered list of two or more realms to which to apply identity rules. When you associate a realm sequence with an identity policy, the Firepower System searches the Active Directory domains in order from first to last as specified in the realm sequence. New/modified screens: **Integration** > **Other Integrations** > **Realms** > **Realm SequencesSystem** > **Integration** > **Realms** > **Realm Sequences**
Realms for user control.	Any	Any	A realm is a connection between the Cloud-Delivered Firewall Management Center either an Active Directory or LDAP user repository.



CHAPTER 80

Identity Source: Captive Portal

- Captive portal identity source, on page 2351
- License requirements for captive portal, on page 2352
- Requirements and prerequisites for captive portal, on page 2352
- Best practices for captive portal guidelines and limitations, on page 2353
- Configure the captive portal for user control, on page 2355
- Troubleshoot the captive portal identity source, on page 2369
- History for captive portal, on page 2370

Captive portal identity source

A captive portal identity source is an active authentication method that

- presents a login page to users and authenticates HTTP and HTTPS traffic to enable secure access to network resources
- requires authentication to access the internet or to access restricted internal resources
- differs from passive authentication by presenting users with a login page on the managed device rather than querying the authentication realm.

You typically use captive portal to require authentication to access the internet or to access restricted internal resources; you can optionally configure guest access to resources. After the system authenticates captive portal users, it handles their user traffic according to access control rules. Captive portal performs authentication on HTTP and HTTPS traffic only.

Captive portal authentication details

Captive portal is one of the authoritative identity sources supported by the system. Active authentication differs from passive authentication in that the user is presented with a login page by the managed device, whereas passive authentication queries the authentication realm (for example, Microsoft AD) to authenticate the user.

After the system authenticates captive portal users, it handles their user traffic according to access control rules. Captive portal performs authentication on HTTP and HTTPS traffic only. You can optionally configure guest access to resources.

Note To use a Microsoft Azure AD (SAML) realm for captive portal, see Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal), on page 2310.

Note HTTPS traffic must be decrypted before captive portal can perform authentication.

Captive portal also records failed authentication attempts. A failed attempt does not add a new user to the list of users in the database. The user activity type for failed authentication activity reported by captive portal is **Failed Auth User**.

The authentication data gained from captive portal can be used for user awareness and user control.

Hostname redirect

A hostname redirect uses the managed device's fully-qualified domain name (FQDN) instead of an IP address for authentication redirections.

(Snort 3 only) An active authentication identity rule redirects to the captive portal port using its configured interface. Because the redirect is typically done to an IP address, the user gets an untrusted certificate error. As this behavior is similar to a man-in-the-middle attack, users might be reluctant to accept the untrusted certificate.

To avoid this problem, you can configure the captive portal to use the managed device's fully-qualified domain name (FQDN). With a properly configured certificate, users will not get an untrusted certificate error, and the authentication will be more seamless and appear to be more secure.

License requirements for captive portal

This reference provides the license requirements for captive portal functionality.

Threat Defense License

Any

Classic License

Control

Requirements and prerequisites for captive portal

This reference provides the requirements and prerequisites for captive portal implementation.

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Best practices for captive portal guidelines and limitations

Captive portal and policies

You configure captive portal in your identity policy and invoke active authentication in your identity rules. Identity policies are associated with access control policies and access control policies define access to resources in the network. For example, you might exclude users in the US-West/Finance group to access Engineering servers or you can prohibit users from accessing nonsecure applications on the network.

You configure some captive portal identity policy settings on the identity policy's **Active Authentication** tab page and configure the rest in the identity rule associated with the access control policy.

An active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected. In each case the system transparently enables or disables TLS/SSL decryption, which restarts the Snort process.



Caution Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include a decryption policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

When the captive portal authenticates users that match an identity rule, any user in a Microsoft Active Directory or LDAP group that has not been downloaded is identified as Unknown. To avoid users being identified as Unknown, configure the realm or realm sequence to download users in all groups you expect to authenticate with captive portal. Unknown users are handled according to the associated access control policy; if the access control policy is configured to block Unknown users, these users are blocked.

To make sure the system downloads all users in a realm or realm sequence, make sure the groups are in the Available Groups list in the realm's configuration.

Microsoft Azure AD SAML realm configuration

To use a Microsoft Azure AD (SAML) realm as a captive portal, see Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal), on page 2302.

Routed interface requirements

Ensure that captive portal active authentication is performed only by a device with a routed interface configured. If you are configuring an identity rule for captive portal and your captive portal device contains inline and routed interfaces, you must configure interface rule conditions in the access control policy to target only the routed interfaces on the device.

If the identity policy associated with your access control policy contains one or more captive portal identity rules and you deploy the policy on the Security Cloud Control that manages one or more devices with routed interfaces configured, the policy deployment succeeds and the routed interfaces perform active authentication.

For more information about synchronizing users and groups, see Synchronize users and groups, on page 2330.

Required certificates and certificate authorities

Before you can use the captive portal for user control and awareness, you must have all of the following:

- To authenticate with Microsoft AD, export the server's root certificate and import it into the Security Cloud Control as a trusted CA certificate.
- An internal certificate object for authenticating with the managed device to which the identity policy is deployed.
- An internal certificate authority for the required decryption rule.

Captive portal requirements and limitations

Note the requirements and limitations when configuring captive portal:

- Captive portal does not support HTTP/3 QUIC connections.
- The system supports up to 20 captive portal logins per second.
- A failed login attempt is counted toward the maximum login attempts only if it occurs within five minutes of the previous failed attempt. The five minute limit is not configurable.

  Connection events show the number of maximum login attempts. To view connection events, click **Analysis** > **Unified Events**.

  If more than five minutes elapse between failed logins, the user is redirected to captive portal for authentication and will not be designated a failed login user or a guest user, and will not be reported to the Security Cloud Control.
- Captive portal does not negotiate TLS v1.0 connections. Only TLS v1.1, v1.2, and TLS 1.3 connections are supported.
- You *cannot* use a realm sequence with the captive portal.
- If a realm is created for a parent domain and the managed device detects a login to a child of that parent domain, the user's subsequent logout is not detected by the managed device.
- To limit the amount of non-HTTP or HTTPS traffic that is allowed through the managed device, you should enter typical HTTP and HTTPS ports in the identity policy's **Ports** tab page.
- Your access control rule must allow traffic destined for the IP address and port of the device you plan to use for captive portal.
- To perform captive portal active authentication on HTTPS traffic, you must use a decryption policy to decrypt the traffic from the users you want to authenticate. You cannot decrypt the traffic in the connection between a captive portal user's web browser and the captive portal daemon on the managed device; this connection is used to authenticate the captive portal user.
- The managed device changes a previously unseen user from **`Pending`** to **`Unknown`** when it determines that the incoming request does not use the HTTP or HTTPS protocol. As soon as the managed device changes a user from **`Pending`** to another state, access control, Quality of Service, and decryption policies

can be applied to that traffic. If your other policies don't permit non-HTTP or HTTPS traffic, configuring ports on the captive portal identity policy can prevent undesired traffic from being allowed through the managed device.

Kerberos prerequisites

If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.

DNS must return a response of 64KB or less to the hostname; otherwise, the AD connection test fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.

Configure the captive portal for user control

Configure captive portal active authentication to enable user authentication through the portal by setting up the necessary components and associating them with access control policies.

To use the captive portal for active authentication, you must set up an LDAP realm; or a Microsoft AD realm or realm sequence; Microsoft Azure AD (SAML) realm; access control policy; an identity policy; a decryption policy; and associate the identity and decryption policies with the same access control policy. Finally, you must deploy the policies to managed devices. This topic provides a high-level summary of those tasks.

Note To use a Microsoft Azure AD (SAML) realm as a captive portal, see Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal), on page 2302.

Before you begin

Perform these tasks first:

- Confirm that your Security Cloud Control manages one or more devices with a *routed* interface configured.
- To use encrypted authentication with the captive portal, either create a PKI object for the authenticating managed device or have your certificate data and key available on the machine from which you're accessing the Security Cloud Control. To create a PKI object, see PKI objects, on page 2986.

Follow these steps to configure the captive portal for user control:

Procedure

Step 1 Create and enable an LDAP realm; or a Microsoft AD realm and optionally realm sequence as discussed in these topics:

- Create an LDAP realm or an Active Directory realm and realm directory, on page 2325
- (Optional.) Create a realm sequence, on page 2332
- Synchronize users and groups, on page 2330

A realm sequence is not supported for the captive portal.

To make sure the system downloads all users in a realm or realm sequence, make sure the groups are in the Available Groups list in the realm's configuration.

For more information, see Synchronize users and groups, on page 2330.

Step 2 Get required certificates and certificate authorities.

You must have all of these:

- To authenticate with Microsoft AD, export the server's root certificate and import it into the Cloud-Delivered Firewall Management Center as a trusted CA certificate.
- An internal certificate object for authenticating with the managed device to which the identity policy is deployed.
- An internal certificate authority for the required decryption rule.

Step 3 Create a network object with an associated trusted certificate authority.

See Create a network object, on page 2357.

Step 4 Create identity policy with an active authentication rule.

The identity policy enables selected users in your realm access resources after authenticating with the captive portal.

For more information, see Create an identity policy and active authentication rule, on page 2358.

Step 5 Configure an access control rule for the captive portal that allows traffic on the captive portal port (by default, TCP 885).

You can choose any available TCP port for the captive portal to use. Whatever your choice, you must create a rule that allows traffic on that port.

For more information, see Create a TCP port access control rule, on page 2360.

Step 6 Add another access control rule to allow users in the selected realm or realm sequence to access resources using the captive portal.

For more information, see Create a user access control rule, on page 2362.

Step 7 Configure a decryption policy with a **Decrypt - Resign** rule for the **Unknown** user so captive portal users can access web pages using the HTTPS protocol.

The captive portal can authenticate users only if the HTTPS traffic is decrypted before the traffic is sent to the captive portal. The captive portal itself is seen by the system as the **Unknown** user.

Create a decryption policy with an outbound rule, on page 2362

Step 8 Associate the identity and decryption policies with the access control policy from step 3.

This final step enables the system to authenticate users with the captive portal.

For more information, see Associate Identity and decryption policies with the access control policy, on page 2366.

What to do next

See Create a network object, on page 2357.

Create a network object

Create a network object to enable access using a fully qualified domain name for captive portal authentication.

This task discusses how to start configuring the captive portal as an identity source.

Before you begin

(Snort 3 only.) Create a fully-qualified host name (FQDN) using your DNS server and upload the Firewall Threat Defense's internal certificate to the Firewall Management Center. You can consult a resource such as this one if you've never done it before. Specify the IP address of a routed interface on one of the devices managed by your Firewall Management Center.

For more information about the network object, see Redirect to host name network rule conditions, on page 2510.

Procedure

Step 1 If you haven't already done so, log in to your Firewall Management Center and click **Objects** > **Object Management**..

Step 2 Expand **PKI**.

Step 3 Click **Internal Certs**.

Step 4 Click **Add Internal Cert**.

a) In the **Name** field, enter a name to identify the internal cert (for example, **`MyCaptivePortal`**).

b) In the **Certificate Data** field, either paste the certificate or use the **Browse** button to locate it.

The certificate Common Name must exactly match the FDQN with which you want captive portal users to authenticate.

c) In the **Key** field, either paste the certificate's private key or use the **Browse** button to locate it.

d) If the certificate is encrypted, select the **Encrypted** check box and enter the password in the adjacent field.

e) Click **Save**.

Step 5 Click **Network**.

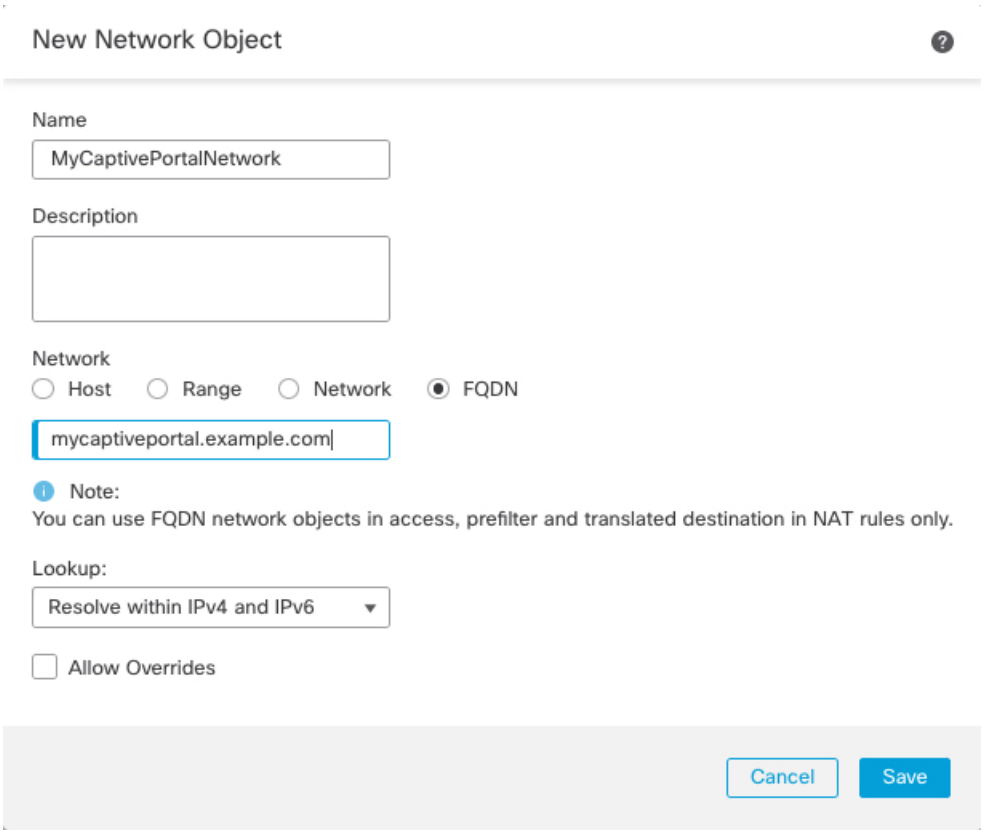
Step 6 Click **Add Network** > **Add Object**.

a) In the **Name** field, enter a name to identify the object (for example, **`MyCaptivePortalNetwork`**).

b) Click **FDQN** and, in the field, enter the name of the captive portal's FDQN.

c) Click an option for **Lookup**.

The following figure shows an example.



d) Click **Save**.

What to do next

Create an identity policy and active authentication rule, on page 2358.

Create an identity policy and active authentication rule

This procedure enables captive portal to perform active authentication for users in your domains using the Secure Firewall Management Center.

This multi-part procedure shows how to set up the captive portal using the default TCP port 885 and using a Cloud-Delivered Firewall Management Center server certificate for both the captive portal and for TLS/SSL decryption. Each part of this example explains one task required to enable the captive portal to perform active authentication.

If you follow all the steps in this procedure, you can configure captive portal to work for users in your domains. You can optionally perform additional tasks, which are discussed in each part of the procedure.

For an overview of the entire procedure, see Configure the captive portal for user control, on page 2355.

Procedure

Step 1 In the Firewall Management Center, click **Policies** > **Access Control heading** > **Identity** and create or edit an identity policy.

Step 2 (Optional.) Click **Add Category** to add a category for the captive portal identity rules and enter a **Name** for the category.

Step 3 Click the **Active Authentication** tab.

a) Choose the appropriate *root* **Server Certificate** from the list or click **Add** (+) to add a certificate.

For more information, see Connect securely to Active Directory or LDAP, on page 2328.

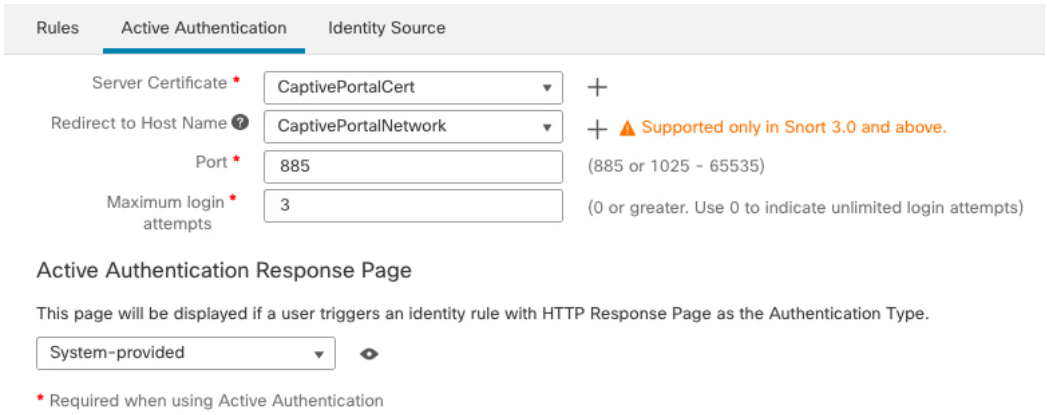
Note
Captive portal does *not* support the use of Digital Signature Algorithm (DSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

b) From the **Redirect to Host Name** field, click the network object you previously created or click **Add** (+).

c) Enter **885** in the **Port** field and specify the **Maximum login attempts**.

d) (Optional.) Choose an **Active Authentication Response Page** as described in Captive portal fields, on page 2367.

This figure shows an example.



Step 4 Click **Save**.

Step 5 Click the **Rules** tab.

Step 6 Click **Add Rule** to add a new captive portal identity policy rule, or click **Edit** (✎) to edit an existing rule.

In the **Add Rule** page,

a) Enter a **Name** for the rule.

b) From the **Action** list, choose **Active Authentication**.

c) Click **Realm & Settings**.

d) From the **Realms** list, choose a realm to use for user authentication

A realm sequence is not supported.

e) (Optional.) Check **Identify as Guest if authentication cannot identify user**. For more information, see Captive portal fields, on page 2367.
f) Choose an **Authentication Protocol** from the list.
g) (Optional.) To exempt specific application traffic from captive portal, see Exclude applications from captive portal, on page 2368.
h) Add conditions to the rule (port, network, and so on) as discussed in Identity rule conditions, on page 2509.
i) Click **Add**.

Step 7 At the top of the page, click **Save**.

What to do next

Continue with Create a TCP port access control rule, on page 2360.

Update a custom authentication form

Allow users to see a list of domains when they authenticate with captive portal by adding a custom realm selection to your authentication form.

After you upgrade to version 7.4.1 (or later) from an earlier release, you must update custom authentication forms for users to see available domains during authentication. This task is always required if you use the HTTP Response Page authentication type; if users authenticate with a realm using another authentication type, this task is optional.

Procedure

Step 1 On the **Active Authentication** tab page of your identity rule, click **Edit** ( ).

Step 2 Enter the following code in the part of the form that requires the user to log in:

Example:

```
<select name="realm" id="realm"></select>
```

Users can now view and select from available domains during captive portal authentication.

Create a TCP port access control rule

Create an access control rule that allows the captive portal to communicate with clients using TCP port 885, enabling secure user authentication.

This part of the procedure shows how to create an access control rule that allows the captive portal to communicate with clients using TCP port 885, which is the captive portal's default port. You can choose another port if you wish, but the port must match the one you chose in Create an identity policy and active authentication rule, on page 2358.

Before you begin

For an overview of the entire captive portal configuration, see Configure the captive portal for user control, on page 2355.

Follow these steps to create a TCP port access control rule:

Procedure

Step 1 Log in to the Firewall Management Center if you have not already done so.

If you haven't done so already, create a certificate for the captive portal as discussed in PKI objects, on page 2986.

Step 2 Click **Policies** > **Access Control heading** > **Access Control** and create or edit an access control policy.

Step 3 Click **Add Rule**.

Step 4 Enter a **Name** for the rule.

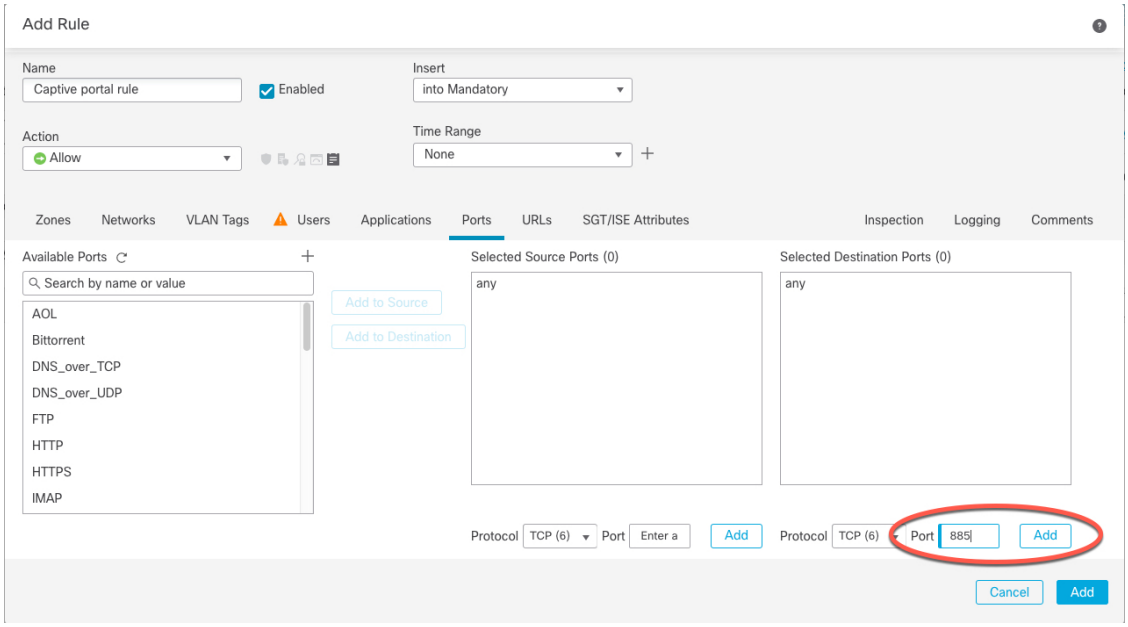
Step 5 Choose **Allow** from the **Action** list.

Step 6 Click **Ports**.

Step 7 From the **Protocol** list under the **Selected Destination Ports** field, choose **TCP**.

Step 8 In the **Port** field, enter **885** and click **Add** next to the **Port** field.

The following figure shows an example.



Step 9 Click **Add** at the bottom of the page.

What to do next

Continue with Create a user access control rule, on page 2362.

Create a user access control rule

Configure an access control rule to allow users in a specified realm to authenticate using captive portal, then adjust rule order in the policy editor to optimize traffic handling and resource use.

This part of the procedure discusses how to add an access control rule that enables users in a realm to authenticate using captive portal.

Before you begin

For an overview of the entire captive portal configuration, see Configure the captive portal for user control, on page 2355.

Procedure

Step 1 In the rule editor, click **Add Rule**.

Step 2 Enter a **Name** for the rule.

Step 3 Choose **Allow** from the **Action** list.

Step 4 Click **Users**.

Step 5 In the **Available Realms** list, click the realms to allow.

If no realms display, click **Refresh** ( ).

Step 6 In the **Available Users** list, choose the users to add to the rule and click **Add to Rule**.

Step 7 (Optional.) Add conditions to the access control policy as discussed in Identity rule conditions, on page 2509.

Step 8 Click **Add**.

Step 9 On the access control rule page, click **Save**.

In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste. Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption

What to do next

Create a decryption policy with an outbound rule, on page 2362

Create a decryption policy with an outbound rule

Create a decryption policy to decrypt and resign outbound traffic before captive portal authentication in the Secure Firewall Management Center.

This part of the procedure discusses how to create a decryption policy to decrypt and resign traffic before the traffic reaches the captive portal. The captive portal can authenticate traffic only after it has been decrypted.

Before you begin

You must have an internal certificate authority (CA) for your outbound server; in other words, the managed device that decrypts the traffic for captive portal users to authenticate. This certificate must be different from the *internal certificate* you use to authenticate the captive portal with the managed device.

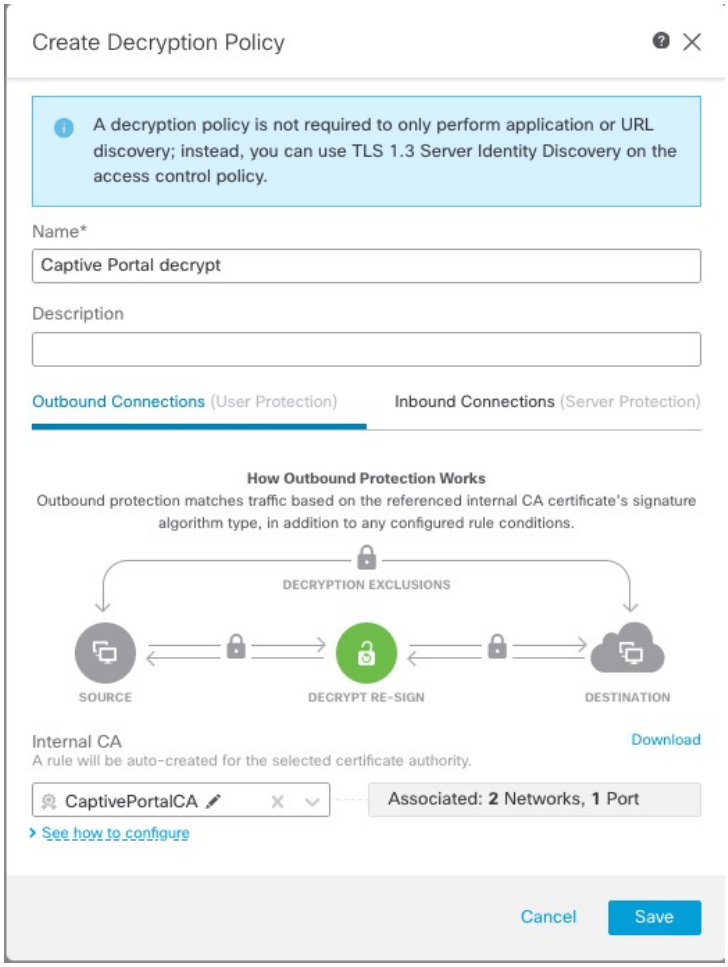
Procedure

Step 1 Click **Policies** > **Access Control heading** > **Decryption**.

Step 2 Click **New Policy**.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

Step 4 Click the **Outbound Connections** tab.



Step 5 Upload or choose certificates for the rules.

The system creates one rule per combination of CA and networks/ports.

Step 6 (Optional.) Choose networks and ports.

For more information:

- Rule-based decryption rule conditions, on page 2201

- Network rule conditions, on page 1701
- Best practice: use application-based rules over port-based rules, on page 1703

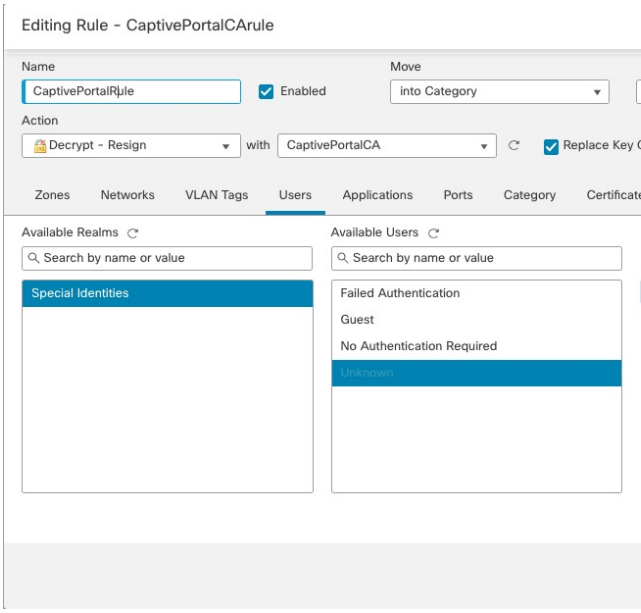
Step 7 Click **Save**.

What to do next

Configure the decryption rule.

Configure a decryption rule

Procedure

	Command or Action	Purpose
Step 1	Click **Edit** ( ) next to the decryption policy you just created.	
Step 2	Click **Edit** ( ) next to the decryption rule for captive portal.	
Step 3	Click **Users**.	
Step 4	Above the **Available Realms** list, click **Refresh** ( ).	
Step 5	In the **Available Realms** list, click **Special Identities**.	
Step 6	In the **Available Users** list, click **Unknown**.	
Step 7	Click **Add to Rule**.	The following figure shows an example. 

	Command or Action	Purpose
Step 8	(Optional.) Set other options as discussed in Rule-based decryption rule conditions, on page 2201.	
Step 9	Click **Add**.	

Create A decryption policy with a Decrypt-Resign rule

Create a decryption policy to decrypt and resign traffic before the traffic reaches the captive portal.

This part of the procedure discusses how to create a decryption policy to decrypt and resign traffic before the traffic reaches the captive portal. The captive portal can authenticate traffic only after it has been decrypted.

Before you begin

For an overview of the entire captive portal configuration, see Configure the captive portal for user control, on page 2355.

Follow these steps to create A decryption policy with a Decrypt-Resign rule:

Procedure

Step 1 If you haven't done so already, log in to the Firewall Management Center and create a certificate object to decrypt TLS/SSL traffic as discussed in PKI objects, on page 2986.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **New Policy**.

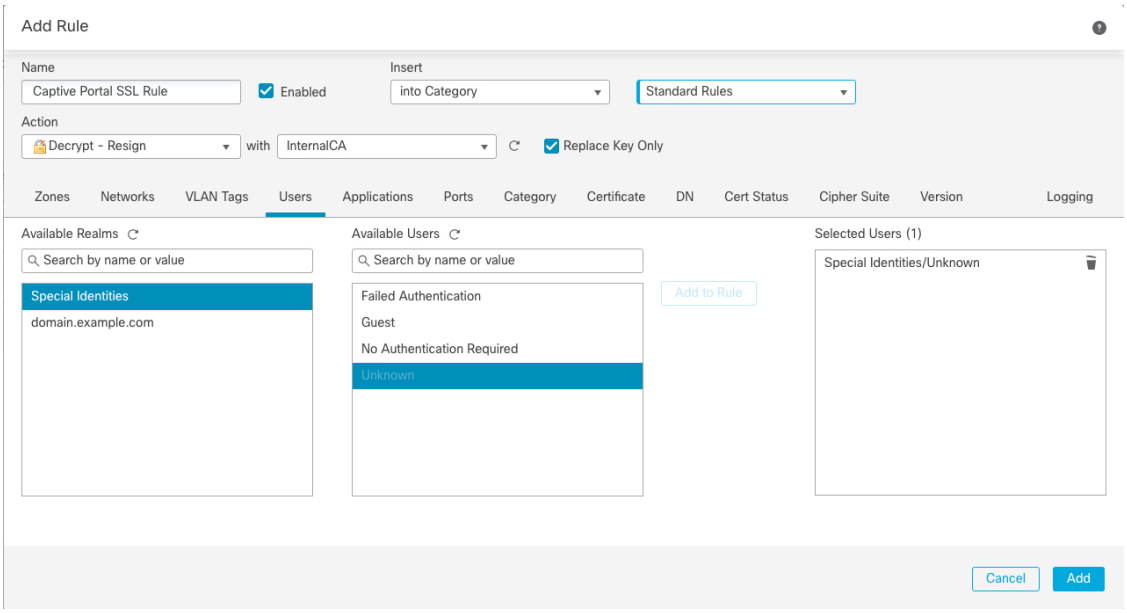
Step 4 Enter a **Name** and choose a **Default Action** for the policy. Default actions are discussed in Rule-based decryption policy default actions, on page 2179.

Step 5 Click **Save**.

Step 6 Click **Add Rule**.

a) Enter a **Name** for the rule.
b) From the **Action** list, choose **Decrypt - Resign**.
c) From the **with** list, choose your PKI object.
d) Under the **Users** tab, refresh the **Available Realms**list by clicking **Refresh** ( ).
e) In the **Available Realms** list, click **Special Identities**.
f) In the **Available Users** list, click **Unknown**.
g) Click **Add to Rule**.

The following figure shows an example.



Step 7 (Optional.) Set other options as discussed in Rule-based decryption rule conditions, on page 2201.

Step 8 Click **Add**.

Step 9 At the top of the page, click **Save**.

What to do next

Associate the identity and decryption policies with the access control policy from step 3.

This final step enables the system to authenticate users with the captive portal.

For more information, see Associate Identity and decryption policies with the access control policy, on page 2366.

Associate Identity and decryption policies with the access control policy

Associate the identity policy and TLS/SSL **Decrypt - Resign** rule with your access control policy to enable captive portal authentication.

This part of the procedure discusses how to associate the identity policy and TLS/SSL **Decrypt - Resign** rule with the access control policy you created earlier. After this, users can authenticate using the captive portal.

For an overview of the entire captive portal configuration, see Configure the captive portal for user control, on page 2355.

Procedure

Step 1 Click **Policies** > **Access Control heading** > **Access Control** and edit the access control policy you created as discussed in Create a TCP port access control rule, on page 2360. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Either create a new access control policy or edit an existing policy.

Step 3 At the top of the page, click the word **Identity**.

Step 4 From the list, choose the name of your identity policy and, at the top of the page, click **Save**.

Step 5 Repeat the preceding steps to associate your captive portal decryption policy with the access control policy.

Step 6 If you haven't done so already, target the policy at managed devices as discussed in Assign access control policy to devices, on page 2014.

What to do next

- Deploy your identity and access control policies to managed devices as discussed in Deploy configuration changes, on page 202.
- Monitor user activity .

Captive portal fields

Use these fields to configure captive portal on the **Active Authentication** tab page of your identity policy. See also Identity rule fields, on page 2517 and Exclude applications from captive portal, on page 2368.

Server Certificate

An internal certificate presented by the captive portal daemon.

Note Captive portal does *not* support the use of Digital Signature Algorithm (DSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

Port

The port number to use for the captive portal connection. You must set up your access control rule with a TCP port to use for the captive portal, then associate the identity policy with that access control policy. For more information, see Create a TCP port access control rule, on page 2360.

Maximum login attempts

The maximum allowed number of failed login attempts before the system denies a user's login request.

Active Authentication Response Page

The system-provided or custom HTTP response page you want to display to captive portal users. After you select an **Active Authentication Response Page** in your identity policy active authentication settings, you also must configure one or more identity rules with **`HTTP Response Page`** as the **Authentication Protocol**.

The system-provided HTTP response page includes **Username** and **Password** fields, a list of realms if you chose to authenticate with a realm sequence, as well as a **Login as guest** button to allow users to access the network as guests. To display a single login method, configure a custom HTTP response page.

An example of what users see when logging in with a response page is shown in Create a sample identity policy with an active authentication rule, on page 2520.

Choose these options:

- To use a generic response, click **System-provided**. You can click **View** ( ) to view the HTML code for this page.
- To create a custom response, click **Custom**. A window with system-provided code is displayed that you can replace or modify. When you are done, save your changes. You can edit a custom page by clicking **Edit** ( ).

Exclude applications from captive portal

This task allows you to exempt specific applications from captive portal active authentication, enabling their traffic to pass through the identity policy without requiring authentication.

You can select applications (identified by their HTTP `User-Agent` strings) and exempt them from captive portal active authentication. This allows traffic from the selected applications to pass through the identity policy without authenticating.

Note Only applications with the **`User-Agent Exclusion` Tag** are displayed in this list.

Before you begin

Follow these steps to exclude applications from captive portal authentication:

Procedure

Step 1 If you haven't done so already, log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity**.

Step 3 Edit the identity policy that contains the captive portal rule.

Step 4 On **Realm & Settings** tab page, expand **HTTP User Agent Exclusions**.

- In the first column, select the check box next to each item to filter applications, then one or more applications, and click **Add to Rule**.

  Check boxes are ANDed together.

- To narrow the filters that are displayed, type a search string in the **Search by name** field; this is especially useful for categories and tags. To clear the search, click **Clear** ( ).
- To refresh the filters list and clear any selected filters, click **Reload** ( ).

Note

The list displays 100 applications at a time.

Step 5 Choose the applications that you want to add to the filter from the **Available Applications** list:

- To narrow the individual applications that appear, enter a search string in the **Search by name** field. To clear the search, click **Clear** ( ).
- Use paging at the bottom of the list to browse the list of individual available applications.
- To refresh the applications list and clear any selected applications, click **Reload** ( ).

Step 6 Add the selected applications to exclude from external authentication. You can click and drag, or you can click **Add to Rule**. The result is the combination of the application filters you selected.

What to do next

- Continue configuring the identity rule as described in Create an identity rule, on page 2515.

Troubleshoot the captive portal identity source

For other related troubleshooting information, see Troubleshoot realms and user downloads, on page 2342 and Troubleshoot user control, on page 2526.

If you experience issues with captive portal, check these items:

- The time on your captive portal managed device must be synchronized with the time on the Cloud-Delivered Firewall Management Center.
- If you have DNS resolution configured and you create an identity rule to perform **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) captive portal, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the hostname you provided when configuring DNS.

  For more information, see Hostname redirect, on page 2352.
- If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.

  DNS must return a response of 64KB or less to the hostname; otherwise, the AD connection test fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.
- If the device you want to use for captive portal contains both inline and routed interfaces, you must configure a zone condition in your captive portal identity rules to target only the routed interfaces on the captive portal device.
- If the captive portal is configured correctly but the redirect to an IP address or fully-qualified domain name (FQDN) fails, disable endpoint security software. This type of software can interfere with the redirection.

- The ERR_CONNECTION_TIMED_OUT error during active authentication occurs when the redirect URL fails to resolve to the managed device's IP address. This error could prevent the client from receiving a user authentication prompt after redirection, cause client traffic on port 885 to be missing from managed device packet captures, or lead to other connection issues. To mitigate connection issues, ensure that the redirect URL resolves to the managed device's IP address on the client's ingress path, where active portal authentication is configured. Verify the configuration by using the `system support captive-portal-debug` command.
- To securely log a user out of captive portal and prevent the user from logging in again, an administrator can close their session: **Analysis** > **Users heading** > **Active Sessions**. If the user closes their browser, the browser does not reauthenticate automatically.
- If you select **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) as the **Authentication Type** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.

  If you select **`HTTP Basic`** as the **Authentication Type** in an identity rule, users on your network might not notice their sessions time out. Most web browsers cache the credentials from **`HTTP Basic`** logins and use the credentials to seamlessly begin a new session after an old session times out.
- If the connection between your Cloud-Delivered Firewall Management Center and a managed device fails, no captive portal logins reported by the device can be identified during the downtime, unless the users were previously seen and downloaded to the Cloud-Delivered Firewall Management Center. The unidentified users are logged as Unknown users on the Cloud-Delivered Firewall Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

  When the captive portal authenticates users that match an identity rule, any user in a Microsoft Active Directory or LDAP group that has not been downloaded is identified as Unknown. To avoid users being identified as Unknown, configure the realm or realm sequence to download users in all groups you expect to authenticate with captive portal. Unknown users are handled according to the associated access control policy; if the access control policy is configured to block Unknown users, these users are blocked.

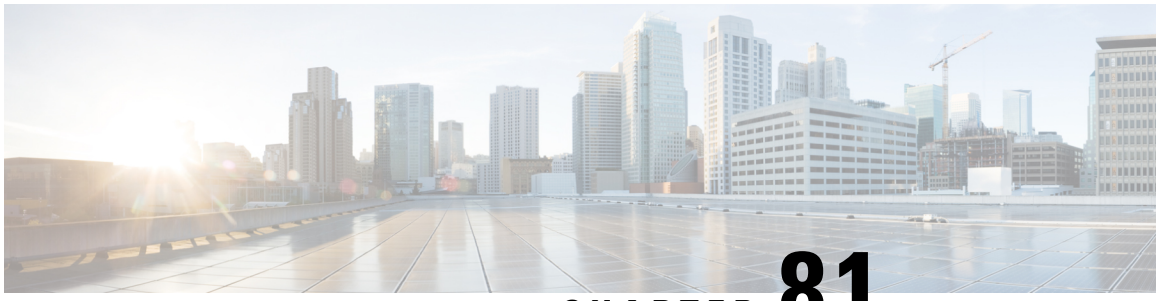
  To make sure the system downloads all users in a realm or realm sequence, make sure the groups are in the Available Groups list in the realm's configuration.

  For more information, see Synchronize users and groups, on page 2330.

History for captive portal

This reference provides the development history and version information for the captive portal feature across different software releases.

Feature	Minimum Cloud-Delivered Firewall Management Center/Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Hostname redirect.	Any	7.1.0 with Snort 3	You can use a network object that contains the fully-qualified host name (FQDN) of the interface that captive portal can use for active authentication requests.
Guest login.	Any	6.1.0	Users can log in as guest using captive portal.
Captive portal.	Any	6.0.0	Feature introduced. You can use the captive portal to require users to enter their credentials when prompted in a browser window. The mapping also allows policies to be based on a user or group of users.



CHAPTER 81

Identity Source: ISE/ISE-PIC

The following topics discuss how to perform user awareness and user control with ISE/ISE-PIC:

- The ISE/ISE-PIC identity source, on page 2373
- License requirements for ISE/ISE-PIC, on page 2375
- Requirements and prerequisites for ISE/ISE-PIC, on page 2375
- Best practices for ISE/ISE-PIC integration, on page 2376
- ISE/ISE-PIC configurations for user control , on page 2379
- ISE/ISE-PIC configuration, on page 2382
- Configure ISE for user control, on page 2387
- Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source, on page 2390
- Troubleshoot Cisco ISE/ISE-PIC or Cisco TrustSec issues, on page 2399
- History for ISE/ISE-PIC, on page 2401

The ISE/ISE-PIC identity source

The ISE/ISE-PIC identity source is an authoritative identity source that

- provides user awareness data for users who authenticate using Active Directory (AD), LDAP, RADIUS, or RSA
- enables passive authentication and user control on Active Directory users
- integrates with the system to manage network access and classify traffic in a Cisco TrustSec environment, and
- supports security group tags (SGT) for access control.

ISE/ISE-PIC integration capabilities

You can integrate your Cisco Identity Services Engine (ISE) or ISE Passive Identity Connector (ISE-PIC) deployment with the system to use ISE/ISE-PIC for passive authentication.

In addition to user awareness and control, if you use ISE Cisco ISE to define and use security group tags (SGT) for classifying traffic in a Cisco TrustSec network, you can write access control rules that use SGT as both source and destination matching criteria. This enables you to block or allow access based on security group membership rather than IP addresses or network objects. Also see Best practices for ISE/ISE-PIC integration, on page 2376.

ISE/ISE-PIC does not report failed login attempts or the activity of ISE Guest Services users.

Note The system does not parse IEEE 802.1x machine authentication but it *does* parse 802.1x user authentication. If you are using 802.1x with ISE, you must include user authentication. 802.1x machine authentication will not provide a user identity to the Cloud-Delivered Firewall Management Center that can be used in policy.

For more information on Cisco ISE/ISE-PIC, see the Identity Services Engine Passive Identity Connector administrator guides or the Identity Services Engine administrator guides.

Note We strongly recommend you use the latest version of ISE/ISE-PIC to get the latest feature set and the most number of issue fixes.

Source and destination security group tag (SGT) matching

If you use Cisco ISE to define and use security group tags (SGT) for classifying traffic in a Cisco TrustSec network, you can write access control rules that use SGT as both source and destination matching criteria. This enables you to block or allow access based on security group membership rather than IP addresses or network objects.

You can also use source SGT matching in DNS rules, as explained in Controlling traffic based on security group tag or dynamic attributes, on page 2572.

Benefits and implementation details

Matching on SGT tags provides these benefits:

- The Cloud-Delivered Firewall Management Center can subscribe to Security Group Tag eXchange Protocol (SXP) mappings from ISE.

  ISE uses SXP to propagate the IP-to-SGT mapping database to managed devices. When you configure Cloud-Delivered Firewall Management Center to use an ISE server, you enable the option to listen to the SXP topic from ISE. This causes the Cloud-Delivered Firewall Management Center to learn about the security group tags and mappings directly from ISE. The Cloud-Delivered Firewall Management Center then publishes SGTs and mappings to managed devices.

  The SXP Topic receives security group tags based on static and dynamic mappings learned through the SXP protocol between ISE and other SXP compliant devices (like switches).

  You can create security group tags in ISE and assign host or network IP addresses to each tag. You can also assign SGTs to user accounts, and the SGT is assigned to the user's traffic. If the switches and routers in the network are configured to do so, these tags then get assigned to packets as they enter the network controlled by ISE, the Cisco TrustSec cloud.

  SXP is *not* supported by ISE-PIC.

- The Cloud-Delivered Firewall Management Center and managed devices can learn about SGT mappings without deploying additional policy. (In other words, you can view connection events for SGT mappings without deploying an access control policy.)

- Supports Cisco TrustSec, which enables you to segment your network to protect critical business assets.

- When a managed device evaluates SGT as a traffic matching criteria for a rule, it uses this priority:
  1. The source SGT tag defined in the packet, if any.

     For the SGT tag to be in the packet, the switches and routers in the network must be configured to add them. See the ISE documentation for information on how to implement this method.

     For the SGT tag to be in the packet, the switches and routers in the network must be configured to add them. See the ISE documentation for information on how to implement this method.
  2. The SGT assigned to the user session, as downloaded from the ISE session directory. The SGT can be matched to source or destination.
  3. The SGT-to-IP address mapping downloaded using SXP. If the IP address is in the range for an SGT, then the traffic matches the rule that uses the SGT. The SGT can be matched to source or destination.

SGT matching examples

- In ISE, create an SGT tag named Guest Users and associate it with the 192.0.2.0/24 network.

  For example, you could use Guest Users as a source SGT condition in your access control rule and restrict access to certain URLs, web site categories, or networks from anyone who accesses your network.
- In ISE, create an SGT tag named Restricted Networks and associate it with the 198.51.100.0/8 network.

  For example, you could use Restricted Networks as a destination SGT rule condition and block access from Guest Users and other networks that have users who are not authorized to access the network.

License requirements for ISE/ISE-PIC

This reference provides license requirements for ISE and ISE-PIC systems.

Threat Defense License

Any

Classic License

Control

Requirements and prerequisites for ISE/ISE-PIC

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Best practices for ISE/ISE-PIC integration

ISE/ISE-PIC version and configuration compatibility

Ensure your ISE/ISE-PIC version and configuration supports proper integration with the Cloud-Delivered Firewall Management Center.

- Use the latest version of ISE/ISE-PIC to get the latest feature set.
- Synchronize the time on the ISE/ISE-PIC server and the Cloud-Delivered Firewall Management Center. Otherwise, the system might perform user timeouts at unexpected intervals.
- Configure and enable a realm for the ISE server assuming the pxGrid persona to implement user control using ISE or ISE-PIC data as described in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
- Ensure each Cloud-Delivered Firewall Management Center host name that connects to an ISE server is unique; otherwise, the connection to one of the Cloud-Delivered Firewall Management Centers will be dropped.
- If you use ISE 3.4 Patch 4 or later as an identity source for in a hybrid deployment, where authorization can use either Microsoft Entra ID (Azure) or on-prem Active Directory, ensure that your ISE authorization policy rule order evaluates on-premises Active Directory external identity source rules before any Entra ID-based authorization rules. This guideline applies to hybrid deployments in which remote access VPN users can authenticate with Microsoft Entra ID, and uses the on-premises Microsoft Active Directory realm for user identity and user-based access control.

For the specific versions of ISE/ISE-PIC that are compatible with this version of the system, see the Cisco Firepower Compatibility Guide.

If you configure ISE/ISE-PIC to monitor a large number of user groups, the system might drop user mappings based on groups due to managed device memory limitations. As a result, rules with realm or user conditions might not perform as expected.

For any device running version 6.7 or later, you can optionally use the **configure identity-subnet-filter** command to limit the subnets that the managed device monitors. For more information, see the *Cisco Secure Firewall Threat Defense Command Reference*.

Alternatively, you can configure a network object and apply that object as an Identity Mapping Filter in the identity policy. See Create an identity policy, on page 2507.

IPv6 support

Verify IPv6 compatibility between ISE/ISE-PIC and your managed devices.

- Compatible versions of ISE/ISE-PIC version 2.x include support for IPv6-enabled endpoints.

- Version 3.0 (patch 2) and later of ISE/ISE-PIC enables IPv6 communication between ISE/ISE-PIC and the Cloud-Delivered Firewall Management Center.

Proxy sequence configuration

A *proxy sequence* is one or more managed devices that can be used to communicate with an LDAP, Active Directory, or ISE/ISE-PIC server. It is necessary only if Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.)

Although you can use one managed device as a proxy sequence, we strongly recommend you set up two or more so that, in the event one managed device cannot communicate with Active Directory or ISE/ISE-PIC, another managed device can take over.

Because of its limited capacity, we recommend you *do not* use a Secure Firewall 220 as a proxy sequence.

Client approval in ISE

Manually approve clients in ISE before the connection between the ISE server and the Cloud-Delivered Firewall Management Center succeeds. Typically, there are two clients: one for the connection test and another for ISE agent.

You can also enable **Automatically approve new accounts** in ISE as discussed in the chapter on Managing users and external identity sources in the *Cisco Identity Services Engine Administrator Guide*.

Unreachable session management

If a user session in ISE/ISE-PIC is reported as unreachable, the system prunes that session so another user with the same IP cannot match the unreachable user's identity rules.

You can control this behavior in ISE/ISE-PIC by going to **Providers** > **Endpoint Probes** and clicking one of these options:

- **Enabled** to cause ISE/ISE-PIC to monitor endpoint connections and therefore to cause the Cloud-Delivered Firewall Management Center prune a session from an unreachable user.
- **Disabled** to cause ISE/ISE-PIC to ignore endpoint connections.

Security group tags (SGT) implementation

A Security Group Tag (SGT) specifies the privileges of a traffic source within a trusted network. Cisco ISE and Cisco TrustSec use a feature called Security Group Access (SGA) to apply SGT attributes to packets as they enter the network. These SGTs correspond to a user's assigned security group within ISE or TrustSec. If you configure ISE as an identity source, the Firepower System can use these SGTs to filter traffic.

Security Group Tags can be used both as source and destination matching criteria in access control and DNS rules.

Note

- To implement user control using only the ISE SGT attribute tag, you do not need to configure a realm for the ISE server. ISE SGT attribute conditions can be configured in policies with or without an associated identity policy.
- In some rules, custom SGT conditions can match traffic tagged with SGT attributes that were *not* assigned by ISE. This is not considered user control, and works only if you are not using ISE/ISE-PIC as an identity source.

To match destination SGT tags in addition to source SGT tags, verify these requirements:

Required ISE version: 2.6 patch 6 or later, 2.7 patch 2 or later

Router support: Any Cisco router that supports SGT inline tagging over Ethernet. For more information, consult a reference such as the Cisco Group Based Policy Platform and Capability Matrix Release

Limitations:

- Quality of Service (QoS) policy uses source SGT matching only; it does *not* use destination SGT matching
- RA-VPN does not receive SGT mappings directly through RADIUS

ISE and high availability

When the primary ISE/ISE-PIC server fails, understand the impact on operations.

As a result of the integration with pxGrid v2, the Cloud-Delivered Firewall Management Center round-robins between both configured ISE hosts until one accepts the connection.

If the connection is lost, the Cloud-Delivered Firewall Management Center resumes round-robin attempts to the connected hosts.

Endpoint location (or location IP)

An Endpoint Location attribute is the IP address of the network device that used ISE to authenticate the user, as identified by ISE.

You must configure and deploy an identity policy to control traffic based on **Endpoint Location (Location IP)**.

ISE attributes

Configuring an ISE connection populates the Cloud-Delivered Firewall Management Center database with ISE attribute data. You can use ISE attributes for user awareness and user control. This is not supported with ISE-PIC.

Endpoint profile (or device type)

An Endpoint Profile attribute is the user's endpoint device type, as identified by ISE.

You must configure and deploy an identity policy to control traffic based on **Endpoint Profile (Device Type)**.

ISE/ISE-PIC configurations for user control

An ISE/ISE-PIC configuration uses realms, identity policies, and access control policies to control user access to network resources.

Configuration options

You can use ISE/ISE-PIC in any of these configurations:

- With a realm, identity policy, and associated access control policy: Use a realm to control *user* access to network resources in policy. You can still use ISE/ISE-PIC Security Group Tags (SGT) metadata in your policies.
- With an access control policy only. No realm or identity policy are necessary: Use this method to control network access using SGT metadata alone.

Configure ISE/ISE-PIC without a realm

Configure ISE to allow or block access to the network using SGT tags for enhanced network security and access control.

This topic provides a high-level overview of tasks you must complete to configure ISE to be able to allow or block access to the network using SGT tags.

Procedure

Step 1 SGT matching: Enable SXP on ISE.

This enables the Cloud-Delivered Firewall Management Center to receive updates from ISE when SGT metadata changes.

Step 2 Export system certificates from ISE/ISE-PIC.

The certificates are required to connect securely between the ISE/ISE-PIC pxGrid, monitoring (MNT) servers and the Cloud-Delivered Firewall Management Center. See Export certificates from the ISE/ISE-PIC server for use in the Cloud-Delivered Firewall Management Center, on page 2384

Step 3 Import the certificates in the Cloud-Delivered Firewall Management Center.

The certificates must be imported as follows:

- pxGrid client certificate: internal certificate with key (**Objects** > **Object Management** > **PKI** > **Internal Certs**)
- pxGrid server certificate: trusted CA (**Objects** > **Object Management** > **PKI** > **Trusted CAs**)
- MNT certificate: trusted CA

Step 4 Create the ISE/ISE-PIC identity source.

The ISE/ISE-PIC identity source enables you to control user activity using Security Group Tags (SGT) provided by ISE/ISE-PIC.

Step 5 Create an access control rule.

The access control rule specifies an action to take (for example, allow or block) if traffic matches the rule criteria. You can use source and destination SGT metadata as matching criteria in the access control rule. See Access control rules, on page 2041.

Step 6 Deploy the access control policy to managed devices.

Before your policy can take effect, it must be deployed to managed devices. See Deploy configuration changes, on page 202.

What to do next

Export certificates from the ISE/ISE-PIC server for use in the Cloud-Delivered Firewall Management Center, on page 2384

Configure ISE/ISE-PIC for user control using a realm

This task enables you to configure ISE/ISE-PIC for user control and to allow or block user or group access to the network. Users and groups can be stored in any server listed in Supported servers for realms, on page 2278.

This topic provides a high-level overview of tasks you must complete to configure ISE/ISE-PIC for user control and to be able to allow or block user or group access to the network.

Procedure

Step 1 Enable SXP on ISE for destination SGT only.

This enables the Cloud-Delivered Firewall Management Center to receive updates from ISE when SGT metadata changes.

Step 2 Export system certificates from ISE/ISE-PIC.

The certificates are required to connect securely between the ISE/ISE-PIC pxGrid, monitoring (MNT) servers and the Cloud-Delivered Firewall Management Center. See the following:

- pxGrid server and MNT server certificate: Export certificates from the ISE/ISE-PIC server for use in the Cloud-Delivered Firewall Management Center, on page 2384
- pxGrid client certificate: Generate a self-signed certificate, on page 2386

Step 3 Import the certificates in the Cloud-Delivered Firewall Management Center.

The certificates must be imported as follows:

- pxGrid client certificate: internal certificate with key (**Objects** > **Object Management** > **PKI** > **Internal Certs**)
- pxGrid server certificate: trusted CA (**Objects** > **Object Management** > **PKI** > **Trusted CAs**)
- MNT certificate: trusted CA

Step 4 Create a proxy sequence for use with the realm and also with ISE/ISE-PIC (optional).

A *proxy sequence* is one or more managed devices that can be used to communicate with an LDAP, Active Directory, or ISE/ISE-PIC server. It is necessary only if Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.)

Although you can use one managed device as a proxy sequence, we strongly recommend you set up two or more so that, in the event one managed device cannot communicate with Active Directory or ISE/ISE-PIC, another managed device can take over.

Step 5 Create a realm.

You must create a realm only to control access to the network by the users and groups you choose.

See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

a) Download users and groups, and enable the realm.

Downloading users and groups enables you to use them in access control rules. See Synchronize users and groups, on page 2330.

Step 6 Create the ISE/ISE-PIC identity source.

The ISE/ISE-PIC identity source enables you to control user activity using Security Group Tags (SGT) provided by ISE/ISE-PIC.

Step 7 Create an identity policy.

An identity policy is a container for one or more identity rules. See Create an identity policy, on page 2507.

a) Create an identity rule.

An identity rule specifies how a realm is used to control access to the network by users and groups. See Create an identity rule, on page 2515.

Step 8 Associate the identity policy with an access control policy.

This enables the access control policy to use users and groups in the realm.

a) Create an access control rule.

The access control rule specifies an action to take (for example, allow or block) if traffic matches the rule criteria. You can use source and destination SGT metadata as matching criteria in the access control rule. See Access control rules, on page 2041.

Step 9 Deploy the access control policy to managed devices.

Before your policy can take effect, it must be deployed to managed devices. See Deploy configuration changes, on page 202.

What to do next

Export certificates from the ISE/ISE-PIC server for use in the Cloud-Delivered Firewall Management Center, on page 2384

ISE/ISE-PIC configuration

This section explains how to configure the ISE/ISE-PIC server for use with identity policies in the Cloud-Delivered Firewall Management Center, including:

- export certificates from the ISE/ISE-PIC server to authenticate with the Cloud-Delivered Firewall Management Center
- publish SXP topics so the Cloud-Delivered Firewall Management Center can be updated with Security Group Tags (SGT) on the ISE server.

Configure security groups and SXP publishing in ISE

This task enables static SGT-to-IP address mappings for Secure Firewall Threat Defense devices, allowing TrustSec policies and access control based on source and destination security group tags (SGT).

There is a lot of configuration that you must do in Cisco Identity Services Engine (ISE) to create the TrustSec policy and security group tags (SGT). Please look at the ISE documentation for more complete information on implementing TrustSec.

The following procedure picks out the highlights of the core settings you must configure in ISE for the Firewall Threat Defense device to be able to download and apply static SGT-to-IP address mappings, which can then be used for source and destination SGT matching in access control rules. See the ISE documentation for detailed information.

The figures in this procedure are based on ISE 2.4. The exact paths to these features might change in subsequent releases, but the concepts and requirements will be the same. Although ISE 2.4 or later is recommended, and preferably 2.6 or later, the configuration should work starting with ISE 2.2 patch 1.

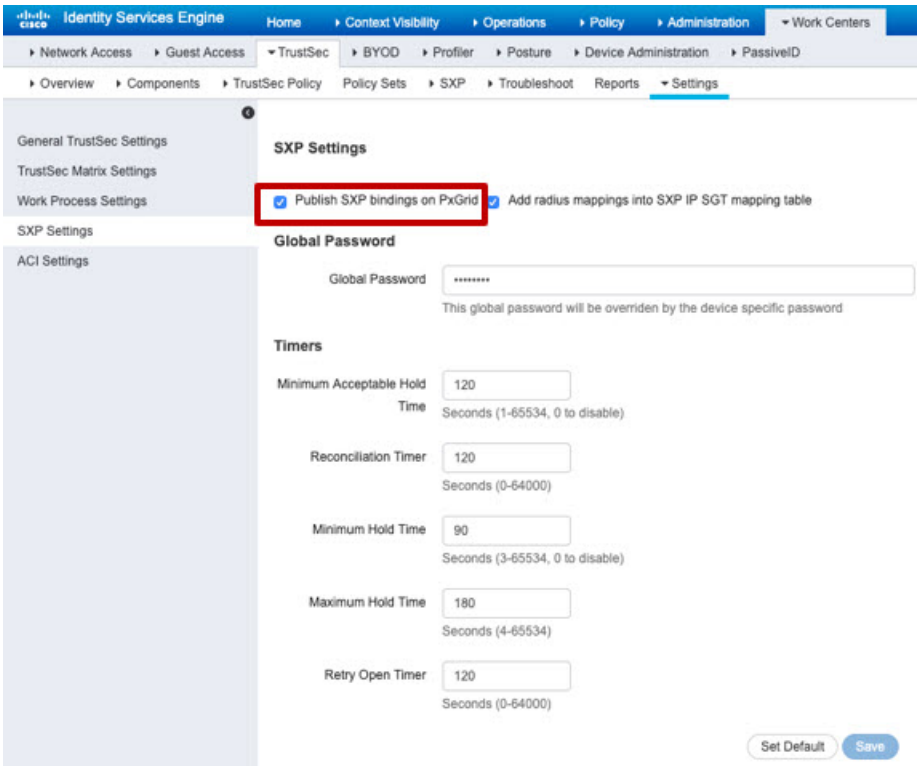
Before you begin

You must have the ISE Plus license to publish SGT-to-IP address static mappings and to get user session-to-SGT mappings so that the Firewall Threat Defense device can receive them.

Procedure

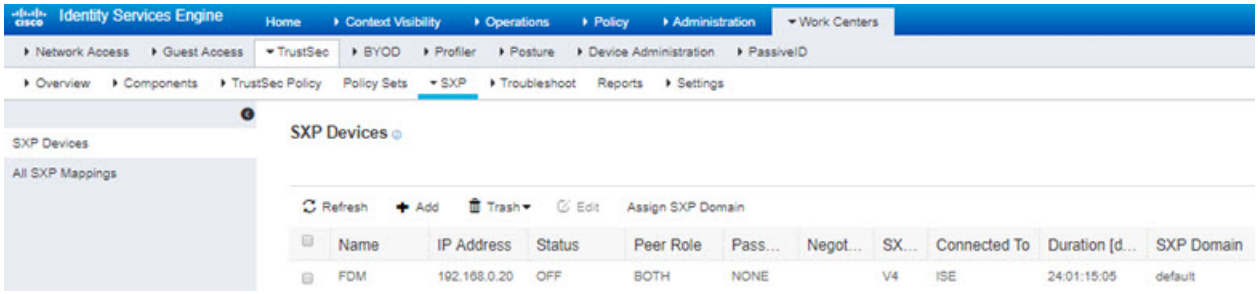
Step 1 Choose **Work Centers** > **TrustSec** > **Settings** > **SXP Settings**, and select the **Publish SXP Bindings on PxGrid** option.

This option makes ISE send the SGT mappings out using SXP. You must select this option for the Firewall Threat Defense device to "hear" anything from listing to the SXP topic. This option must be selected for the Firewall Threat Defense device to get static SGT-to-IP address mapping information. It is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

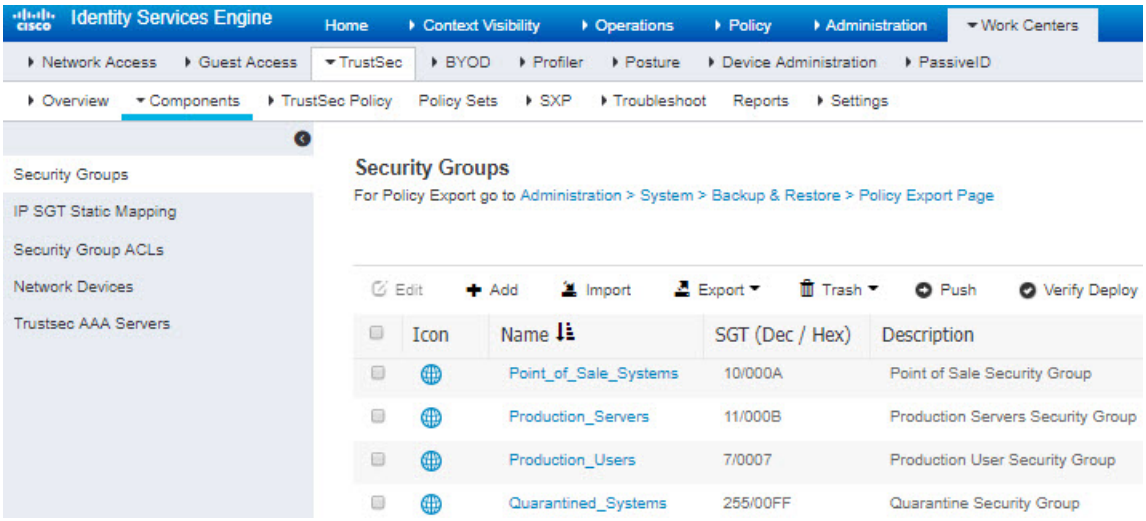


Step 2 Choose **Work Centers** > **TrustSec** > **SXP** > **SXP Devices**, and add a device.

This does not have to be a real device, you can even use the management IP address of the Firewall Threat Defense device. The table simply needs at least one device to induce ISE to publish the static SGT-to-IP address mappings. This step is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

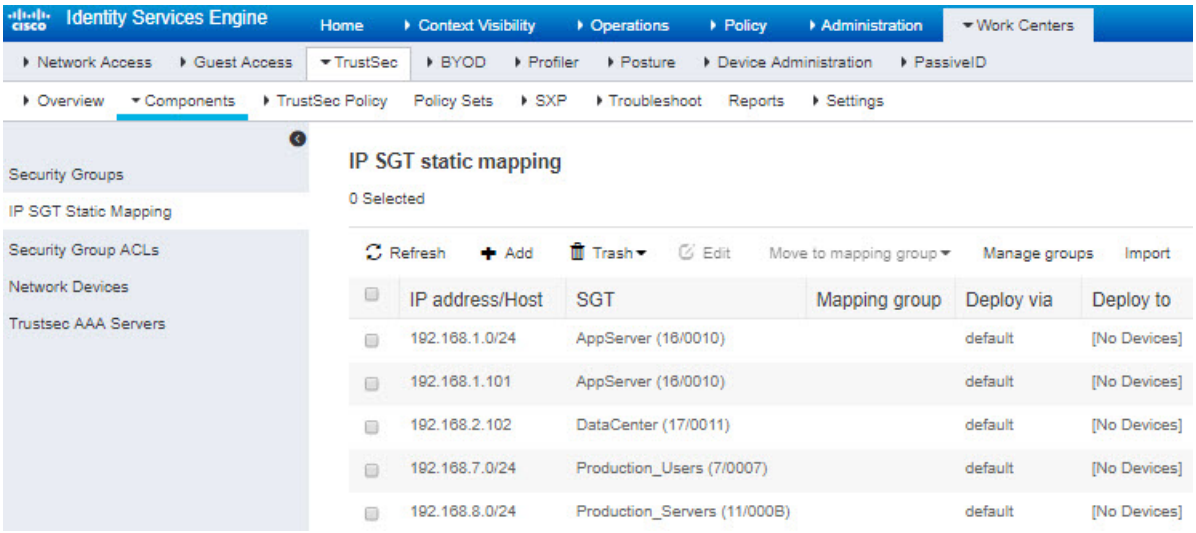


Step 3 Choose **Work Centers** > **TrustSec** > **Components** > **Security Groups** and verify there are security group tags defined. Create new ones as necessary.



Step 4 Choose **Work Centers** > **TrustSec** > **Components** > **IP SGT Static Mapping** and map host and network IP addresses to the security group tags.

This step is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.



You have configured ISE with the necessary security groups, SXP publishing settings, and static SGT-to-IP mappings required for Secure Firewall Threat Defense devices to receive and apply TrustSec policies.

Export certificates from the ISE/ISE-PIC server for use in the Cloud-Delivered Firewall Management Center

Certificate exported from Identity Services Engine (ISE) / Identity Services Engine - Passive Identity Connector (ISE-PIC) servers enables secure authentication between the Cloud-Delivered Firewall Management Center and ISE systems.

You might need to export one, or as many as three, certificates, depending on how your ISE system is set up:

- One certificate for the pxGrid server
- One certificate for the monitoring (MNT) server
- One certificate, including the private key, for the pxGrid client (that is, the Cloud-Delivered Firewall Management Center)

  Unlike the first two certificates, this is a self-signed certificate.

These certificates are then imported into the Cloud-Delivered Firewall Management Center:

- pxGrid client certificate: internal certificate with key (**Objects** > **Object Management** > **PKI** > **Internal Certs**)
- pxGrid server certificate: trusted CA (**Objects** > **Object Management** > **PKI** > **Trusted CAs**)
- MNT certificate: trusted CA

Export a system certificate

Export a system certificate or a certificate with its associated private key for backup or transfer to another system.

You can export a system certificate or a certificate and its associated private key. If you export a certificate and its private key for backup purposes, you can reimport them later if needed.

Before you begin

To perform the following task, you must be a Super Admin or System Admin.

Procedure

Step 1 In the Cisco ISE GUI, click the **Menu** icon (☰) and choose **Administration** > **System** > **Certificates** > **System Certificates**.

Step 2 Check the check box next to the certificate that you want to export and click **Export**.

Step 3 Choose whether to export only the certificate, or the certificate and its associated private key.

Tip
We do not recommend exporting the private key that is associated with a certificate because its value may be exposed. If you must export a private key (for example, when you export a wildcard system certificate to be imported into the other Cisco ISE nodes for inter-node communication), specify an encryption password for the private key. You must specify this password while importing this certificate into another Cisco ISE node to decrypt the private key.

Step 4 Enter the password if you have chosen to export the private key. The password should be at least eight characters long.

Step 5 Click **Export** to save the certificate to the file system that is running your client browser.

If you export only the certificate, the certificate is stored in the PEM format. If you export both the certificate and private key, the certificate is exported as a .zip file that contains the certificate in the PEM format and the encrypted private key file.

Generate a self-signed certificate

Generate a self-signed certificate to add a new local certificate for internal testing and evaluation purposes in Cisco ISE.

Add a new local certificate by generating a self-signed certificate. Cisco recommends that you only employ self-signed certificates for your internal testing and evaluation needs. If you plan to deploy Cisco ISE in a production environment, use CA-signed certificates whenever possible to ensure more uniform acceptance around a production network.



Note If you use a self-signed certificate and you want to change the hostname of your Cisco ISE node, log in to the administration portal of the Cisco ISE node, delete the self-signed certificate that has the old hostname, and generate a new self-signed certificate. Otherwise, Cisco ISE continues to use the self-signed certificate with the old hostname.

Before you begin

To perform this task, you must be a Super Admin or System Admin.

Follow these steps to generate a self-signed certificate:

Procedure

Step 1 In the Cisco ISE GUI, click the **Menu** icon (☰) and choose **Administration** > **System** > **Certificates** > **System Certificates**.

To generate a self-signed certificate from a secondary node, choose **Administration** > **System** > **Server Certificate**.

Step 2 In the ISE-PIC GUI, click the **Menu** icon (☰) and choose **Certificates** > **System Certificates**.

Step 3 Click **Generate Self Signed Certificate** and enter the details in the window displayed.

Step 4 Check the check boxes in the **Usage** area based on the service for which you want to use this certificate.

Step 5 Click **Submit** to generate the certificate.

To restart the secondary nodes, from the CLI, enter these commands in this order:

a) **application stop ISE**
b) **application start ISE**

Import ISE/ISE-PIC certificates

Import certificates and keys from your ISE or ISE-PIC server to the firewall to enable secure pxGrid and MNT integration for cloud-delivered firewall management.

This procedure is optional.

- pxGrid client certificate: internal certificate with key (**Objects** > **Object Management** > **PKI** > **Internal Certs**)
- pxGrid server certificate: trusted CA (**Objects** > **Object Management** > **PKI** > **Trusted CAs**)
- MNT certificate: trusted CA

Before you begin

Export certificates from the ISE/ISE-PIC server as discussed in Export a system certificate, on page 2385. The certificates and key must be present on the machine from which you log in to the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Log in to the Cloud-Delivered Firewall Management Center if you have not already done so.
Step 2 Click **Objects** > **Object Management** > **PKI** > **Internal Certs**.
Step 3 Click **Add Internal Cert**.
Step 4 Follow the prompts on your screen to import the certificate and private key.
Step 5 Click **Trusted CAs**.
Step 6 Click **Add Trusted CA**.
Step 7 Follow the prompts on your screen to import the pxGrid server certificate.
Step 8 Repeat the preceding steps, if necessary, to import the MNT server's trusted CA.

Configure ISE for user control

Configure ISE/ISE-PIC identity source to enable user session retrieval from Microsoft Active Directory or LDAP servers for identity-based access control.

The following procedure discusses how to configure the ISE/ISE-PIC identity source. You must be in the global domain to perform this task.

Firewall Threat Defense Feature History:

7.2—Optionally add a proxy, which is a connection to one or more cSecurity Cloud Control in the event Security Cloud Control cannot communicate with the ISE/ISE-PIC server. .

Before you begin

- To get user sessions from a Microsoft Active Directory Server or supported LDAP server, configure and enable a realm for the Cisco ISE server, assuming the pxGrid persona, as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

- To get all mappings that are defined in Cisco ISE, including SGT-to-IP address mappings published through SXP, use the procedure that follows. As an alternative, you have the following options:
  - To use the SGT information in the packets only, and not use mappings downloaded from Cisco ISE, skip the steps discussed in Create and edit access control rules, on page 2066. Note that in this case, you can use SGT tags as a source condition only; these tags will never match destination criteria.
  - To use SGT in packets and user-to-IP-address/SGT mappings only, do not subscribe to the SXP topic in the Cisco ISE identity source, and do not configure ISE to publish SXP mappings. You can use this information for both source and destination matching conditions.
- (Advanced configuration only.) Export certificates from the ISE/ISE-PIC server and optionally import them into the Cloud-Delivered Firewall Management Center as discussed in Export certificates from the ISE/ISE-PIC server for use in the Cloud-Delivered Firewall Management Center, on page 2384.
- To publish SXP topics so the Cloud-Delivered Firewall Management Center can be updated with Security Group Tags (SGT) on the ISE server, see ISE/ISE-PIC configuration, on page 2382.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Integration** > **Other Integrations** > **Identity Sources**.

Step 3 Click **Identity Services Engine** for the **Service Type** to enable the ISE connection.

Note
To disable the connection, click **None**.

Step 4 Enter a **Primary Host Name/IP Address** and, optionally, a **Secondary Host Name/IP Address**.

Step 5 Click the appropriate certificate authorities from the **pxGrid Server CA** and **MNT Server CA** lists, and the appropriate certificate from the **pxGrid Client Certificate** list. You can also click **Add** (+) to add a certificate.

Note
The **pxGrid Client Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.

Step 6 (Optional.) Enter an **ISE Network Filter** using CIDR block notation.

Step 7 In the Subscribe To section, check the following:

- **Session Directory Topic** to receive ISE user session information from the ISE server.
- **SXP Topic** to receive updates to SGT-to-IP mappings when available from the ISE server. This option is required to use destination SGT tagging in access control rules.

Step 8 (Optional.) From the **Proxy** list, click either a managed device or a proxy sequence.
If Security Cloud Control cannot communicate with your ISE/ISE-PIC server, you can choose either a managed device or proxy sequence to do it. For example, your Security Cloud Control might be in a public cloud but the ISE/ISE-PIC server might be on an internal intranet.

Step 9 To test the connection, click **Test**.

If the test fails, click **Additional Logs** for more information about the connection failure.

Note
If you set up load balancing with ISE, the node that responds to the test should be the currently active node. For more information, see Cisco pxGrid Node.

What to do next

- Specify users to control and other options using an identity policy as described in Create an identity policy, on page 2507.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Use Security Group Tags (SGT) from Cisco ISE as dynamic attributes in access control policies.
- Deploy your identity and access control policies to managed devices as discussed in Deploy configuration changes, on page 202.
- Monitor user activity .

ISE/ISE-PIC configuration fields

The fields in this reference are used to configure a connection to ISE/ISE-PIC.

Connection and certificate configuration fields for ISE/ISE-PIC integration:

- **Primary and Secondary Host Name/IP Address**: The hostname or IP address for the primary and, optionally, the secondary pxGrid ISE servers. The ports used by the host names you specify must be reachable by both ISE and the Cloud-Delivered Firewall Management Center.
- **pxGrid Server CA**: The trusted certificate authority for the pxGrid framework. If your deployment includes a primary and a secondary pxGrid node, the certificates for both nodes must be signed by the same certificate authority.
- **MNT Server CA**: The trusted certificate authority for the ISE certificate when performing bulk downloads. If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.
- pxGrid Client Certificate: The internal certificate and key that the Cloud-Delivered Firewall Management Center must provide to /ISE-PIC to connect to /ISE-PIC or to perform bulk downloads.

Note The **pxGrid Client Certificate** must include the clientAuth extended key usage value, or it must not include any extended key usage values.

Data filtering, subscription, and proxy configuration fields:

- **ISE Network Filter**: An optional filter you can set to restrict the data that ISE reports to the Cloud-Delivered Firewall Management Center. If you provide a network filter, ISE reports data from the networks within that filter. You can specify a filter in these ways:
  - Leave the field blank to specify **`any`**.

- Enter a single IPv4 address block using CIDR notation.
- Enter a list of IPv4 address blocks using CIDR notation, separated by commas.

- **Subscribe to**: **Session Directory Topic**: Check this box to subscribe to user session information from the ISE server. Includes SGT and endpoint metadata. **SXP Topic**: Check this box to subscribe to SXP mappings from the ISE server.
- **Proxy**: You can optionally choose either a managed device or a proxy sequence to communicate with ISE/ISE-PIC if Security Cloud Control is unable to do so. For example, your Security Cloud Control might be in a public cloud but the ISE/ISE-PIC server might be on an internal intranet.



Note This version of the system does not support filtering using IPv6 addresses, regardless of your ISE version.

Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source

Cisco ISE identity source configuration methods are approaches that

- enable authentication and certificate management for the firewall
- provide quick or advanced setup options based on integration requirements, and
- support different ISE deployment types including ISE-PIC.

Configuration methods

You can configure the Cisco ISE identity source in these ways:

- Quick configuration (new): *Cisco ISE only; not supported for ISE-PIC*. Enter the user name and password of a user in at least the External RESTful Services (ERS) Operator group. The Cloud-Delivered Firewall Management Center logs into the Cisco ISE Primary Authentication Node (PAN), downloads the certificates, and configures the identity source. You have the option of saving the ISE configuration for later use.

  For more information about groups, see the section on Cisco ISE Administrator Groups in the Identity Services Engine administrator guides.

  For details, see Quick configuration, on page 2392.

- Advanced configuration (old): Same as previous Cloud-Delivered Firewall Management Center releases and works for both Cisco ISE and ISE-PIC. You must get certificates and other information from your ISE/ISE-PIC server as discussed in ISE/ISE-PIC configuration, on page 2382.

  For details about advanced configuration, see Configure Cisco ISE advanced configuration, on page 2396.

Cisco ISE quick configuration

Cisco ISE quick configuration lets you configure ISE using only the user name and password of a user in the External RESTful Services (ERS) Operator group. (This feature is available for ISE only, not ISE-PIC.)

Prerequisites and requirements

Complete these tasks first: Configure prerequisites for Cisco ISE quick configuration, on page 2391.

You must have the following information:

- the fully qualified domain name or IP address of the policy administration node (PAN) and
- the user name and password of a user in at least the External RESTful Services (ERS) Operator group.

Note the following about SGT-to-IP mappings and SXP:

- To get all mappings that are defined in Cisco ISE, including SGT-to-IP address mappings published through SXP, use the procedure that follows. As an alternative, you have the following options:
  - To use the SGT information in the packets only, and not use mappings downloaded from Cisco ISE, skip the steps discussed in Create and edit access control rules, on page 2066. Note that in this case, you can use SGT tags as a source condition only; these tags will never match destination criteria.
  - To use SGT in packets and user-to-IP-address/SGT mappings only, do not subscribe to the SXP topic in the Cisco ISE identity source, and do not configure ISE to publish SXP mappings. You can use this information for both source and destination matching conditions.

Configure prerequisites for Cisco ISE quick configuration

Configure the necessary ISE settings to enable quick configuration functionality through API services and certificate authority features.

As an ISE administrator, you must configure specific API and certificate authority settings to ensure ISE quick configuration works properly.

Before you begin

You must have ISE administrator privileges.

Follow these steps to configure prerequisites for Cisco ISE quick configuration:

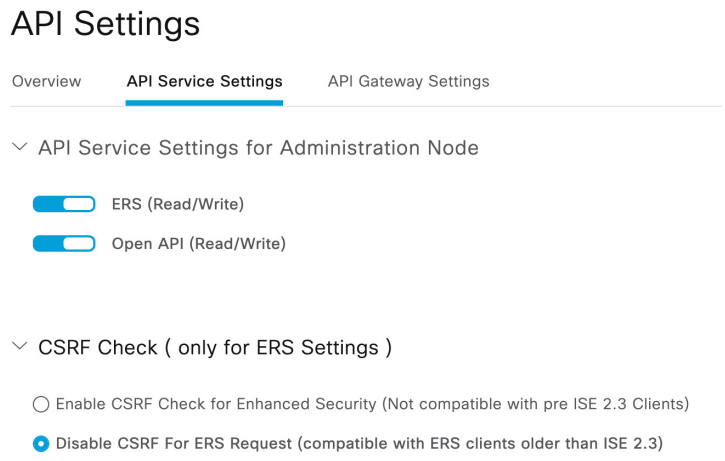
Procedure

Step 1 Enable External RESTful Services (ERS) (read/write), Open API (read/write), and *disable* the Cross-Site Request Forgery (CSRF) check.

Go to  > **Administration** > **System** > **Settings** > **API Settings** > **API Service Settings**

For more information, see Enable API Service in the Identity Services Engine administrator guides.

The following figure shows an example.



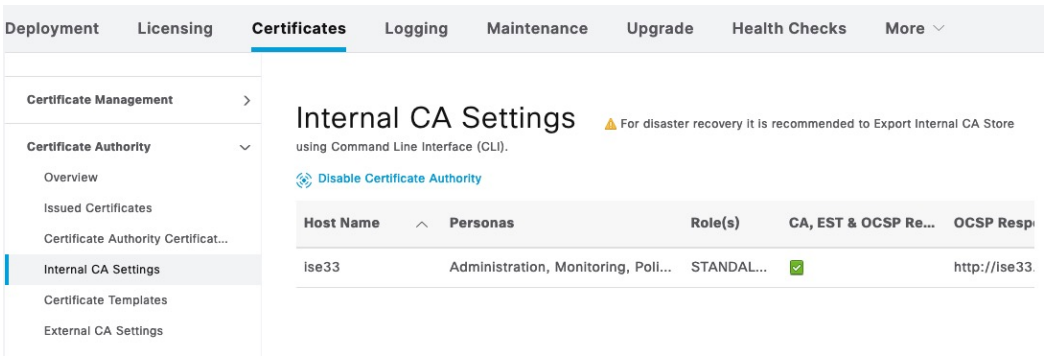
Step 2 Slide **ERS (Read/Write)** to **Enabled**.

Step 3 Enable the certificate authority.

Go to > **Administration** > **System** > **Certificates** > **Certificate Authority** > **Internal CA Settings**

For more information, see Internal CA Settings in the Identity Services Engine administrator guides.

The following figure shows an example.



ISE quick configuration is now enabled with the proper API services and certificate authority settings configured.

Quick configuration

Configure Cisco ISE integration with the firewall to authenticate applications and enable identity-based access control.

This task discusses how to configure Cisco ISE (but not ISE-PIC) by entering a user name and password. The Cloud-Delivered Firewall Management Center then logs in to ISE and downloads the necessary certificates to authenticate the two applications.

Firewall Threat Defense Feature History:

7.6—This feature is introduced.

Before you begin

See these topics first:

- Cisco ISE quick configuration, on page 2391
- Configure prerequisites for Cisco ISE quick configuration, on page 2391

Procedure

Step 1 Log in to the Cloud-Delivered Firewall Management Center. Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**

Step 2 Click **Identity Services Engine** for the **Service Type** to enable the ISE connection.

Note
To disable the connection, click **None**.

Step 3 Click **Quick Configuration (New)**.

a) In the **Primary PAN FQDN/IP Address** field, enter the fully qualified domain name or IP address of the policy administration node (PAN). Do *not* enter a scheme (such as **https://**).

b) In the **Username** field, enter the user name of a user in at least the ERS Operator group.

For more information about groups, see the section on Cisco ISE Administrator Groups in the Identity Services Engine administrator guides.

c) In the **Password** field, enter the user's password.

Step 4 (Optional.) Enter an **ISE Network Filter** using CIDR block notation.

Step 5 In the Subscribe To section, check both options:

- **Session Directory Topic** to receive ISE user session information from the ISE server.
- **SXP Topic** to receive updates to SGT-to-IP mappings when available from the ISE server. This option is required to use destination SGT tagging in access control rules.

Step 6 (Optional.) From the **Proxy** list, click either a managed device or a proxy sequence.
If Security Cloud Control cannot communicate with your ISE/ISE-PIC server, you can choose either a managed device or proxy sequence to do it. For example, your Security Cloud Control might be in a public cloud but the ISE/ISE-PIC server might be on an internal intranet.

Step 7 To test the connection, click **Test**.

Step 8 (Optional.) After a successful test, click **Save this Config** at the top of the page to save the configuration on the Cloud-Delivered Firewall Management Center.

What to do next

- Specify users to control and other options using an identity policy as described in Create an identity policy, on page 2507.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

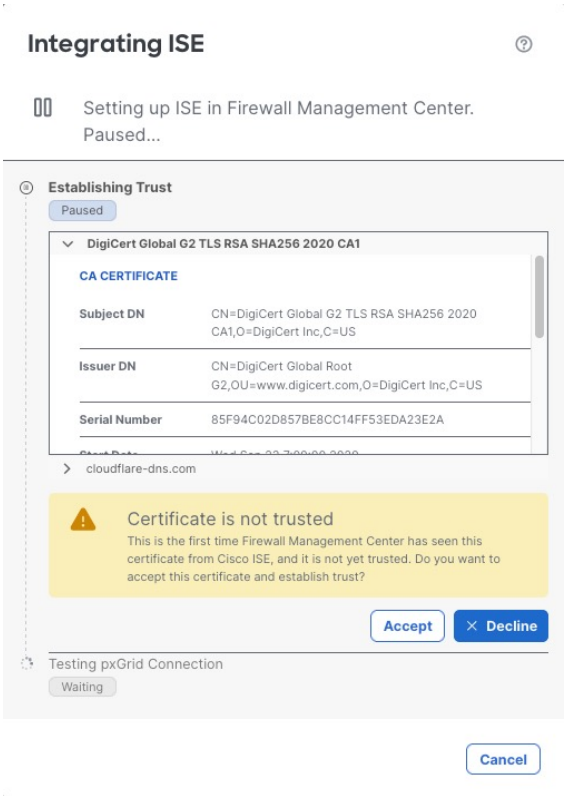
- Use Security Group Tags (SGT) from Cisco ISE as dynamic attributes in access control policies.
- Deploy your identity and access control policies to managed devices as discussed in Deploy configuration changes, on page 202.
- Monitor user activity .

Cisco Identity Services Engine (Cisco ISE) quick configuration results

This reference provides information about the different results that can occur during Cisco Identity Services Engine (Cisco ISE) quick configuration integration with the firewall, including certificate trust prompts and authentication outcomes.

Initial information page for an untrusted certificate

When you enter required information in the **Quick Configuration (New)** tab page, the system first checks to see if the certificate was signed by a trusted authority. If not, a page similar to this is displayed.



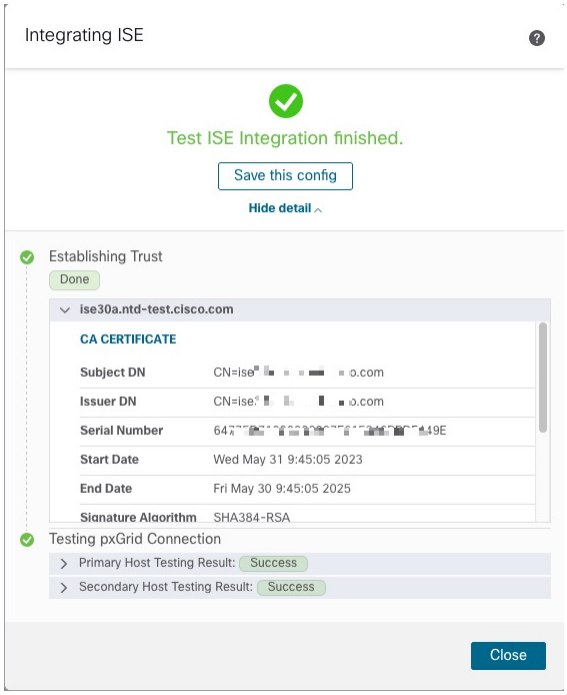
You have these options:

- Click **Accept** for the > to trust the certificate and continue authenticating with ISE.
- Click **Decline** to quit without authenticating with ISE.

  You can then verify the ISE configuration on the **Quick Configuration (New)** tab page and try again.

Successful ISE integration

If the Cloud-Delivered Firewall Management Center was able to authenticate successfully with the information you provided, a page similar to this is displayed.

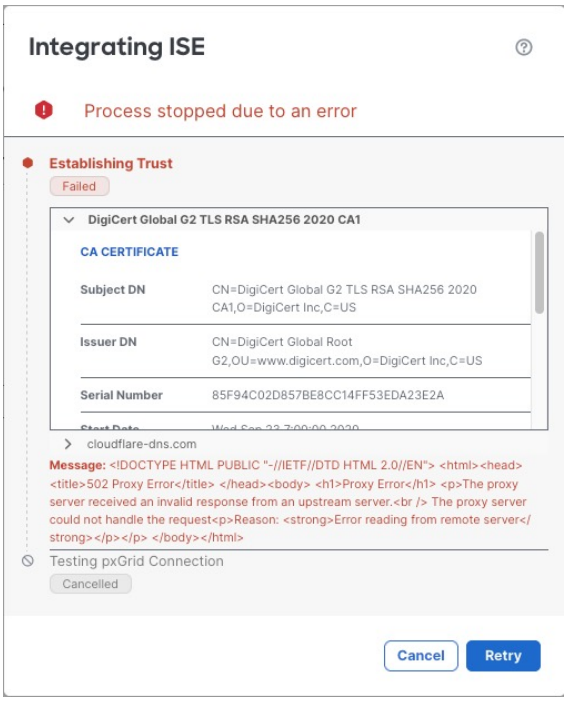


Click **Save this Config** to save it and be able to use the identity source in the future.

Optionally expand one of the lists under Testing pxGrid Connection to display certificates downloaded from those sources.

Unsuccessful ISE integration

If the Cloud-Delivered Firewall Management Center fails to authenticate with ISE for any reason, a page similar to this is displayed.



The error messages are displayed in red text.

You have these options:

- If the error is temporary (for example, a temporarily network issue), click **Retry**.
- To try again with different ISE credentials, click **Cancel**.

Configure Cisco ISE advanced configuration

This task configures ISE/ISE-PIC identity source integration with the firewall to enable advanced security policies based on user identity.

The following procedure discusses how to configure the ISE/ISE-PIC identity source. You must be in the global domain to perform this task.

Firewall Threat Defense Feature History:

7.2—Optionally add a proxy, which is a connection to one or more cSecurity Cloud Control in the event Security Cloud Control cannot communicate with the ISE/ISE-PIC server. .

Before you begin

- To get user sessions from a Microsoft Active Directory Server or supported LDAP server, configure and enable a realm for the Cisco ISE server, assuming the pxGrid persona, as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
- To get all mappings that are defined in Cisco ISE, including SGT-to-IP address mappings published through SXP, use the procedure that follows. As an alternative, you have the following options:

- To use the SGT information in the packets only, and not use mappings downloaded from Cisco ISE, skip the steps discussed in Create and edit access control rules, on page 2066. Note that in this case, you can use SGT tags as a source condition only; these tags will never match destination criteria.
- To use SGT in packets and user-to-IP-address/SGT mappings only, do not subscribe to the SXP topic in the Cisco ISE identity source, and do not configure ISE to publish SXP mappings. You can use this information for both source and destination matching conditions.

- (Advanced configuration only.) Export certificates from the ISE/ISE-PIC server and optionally import them into the Cloud-Delivered Firewall Management Center as discussed in Export certificates from the ISE/ISE-PIC server for use in the Cloud-Delivered Firewall Management Center, on page 2384.
- To publish SXP topics so the Cloud-Delivered Firewall Management Center can be updated with Security Group Tags (SGT) on the ISE server, see ISE/ISE-PIC configuration, on page 2382.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Integration** > **Other Integrations** > **Identity Sources**.

Step 3 Click **Identity Services Engine** for the **Service Type** to enable the ISE connection.

Note
To disable the connection, click **None**.

Step 4 Enter a **Primary Host Name/IP Address** and, optionally, a **Secondary Host Name/IP Address**.

Step 5 Click the appropriate certificate authorities from the **pxGrid Server CA** and **MNT Server CA** lists, and the appropriate certificate from the **pxGrid Client Certificate** list. You can also click **Add** (+) to add a certificate.

Note
The **pxGrid Client Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.

Step 6 (Optional.) Enter an **ISE Network Filter** using CIDR block notation.

Step 7 In the Subscribe To section, check the following:

- **Session Directory Topic** to receive ISE user session information from the ISE server.
- **SXP Topic** to receive updates to SGT-to-IP mappings when available from the ISE server. This option is required to use destination SGT tagging in access control rules.

Step 8 (Optional.) From the **Proxy** list, click either a managed device or a proxy sequence.
If Security Cloud Control cannot communicate with your ISE/ISE-PIC server, you can choose either a managed device or proxy sequence to do it. For example, your Security Cloud Control might be in a public cloud but the ISE/ISE-PIC server might be on an internal intranet.

Step 9 To test the connection, click **Test**.

If the test fails, click **Additional Logs** for more information about the connection failure.

Note

If you set up load balancing with ISE, the node that responds to the test should be the currently active node. For more information, see Cisco pxGrid Node.

What to do next

- Specify users to control and other options using an identity policy as described in Create an identity policy, on page 2507.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Use Security Group Tags (SGT) from Cisco ISE as dynamic attributes in access control policies.
- Deploy your identity and access control policies to managed devices as discussed in Deploy configuration changes, on page 202.
- Monitor user activity .

ISE/ISE-PIC configuration fields

The fields in this reference are used to configure a connection to ISE/ISE-PIC.

Connection and certificate configuration fields for ISE/ISE-PIC integration:

- **Primary and Secondary Host Name/IP Address**: The hostname or IP address for the primary and, optionally, the secondary pxGrid ISE servers. The ports used by the host names you specify must be reachable by both ISE and the Cloud-Delivered Firewall Management Center.
- **pxGrid Server CA**: The trusted certificate authority for the pxGrid framework. If your deployment includes a primary and a secondary pxGrid node, the certificates for both nodes must be signed by the same certificate authority.
- **MNT Server CA**: The trusted certificate authority for the ISE certificate when performing bulk downloads. If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.
- pxGrid Client Certificate: The internal certificate and key that the Cloud-Delivered Firewall Management Center must provide to /ISE-PIC to connect to /ISE-PIC or to perform bulk downloads.

Note The **pxGrid Client Certificate** must include the clientAuth extended key usage value, or it must not include any extended key usage values.

Data filtering, subscription, and proxy configuration fields:

- **ISE Network Filter**: An optional filter you can set to restrict the data that ISE reports to the Cloud-Delivered Firewall Management Center. If you provide a network filter, ISE reports data from the networks within that filter. You can specify a filter in these ways:
  - Leave the field blank to specify **`any`**.
  - Enter a single IPv4 address block using CIDR notation.

- Enter a list of IPv4 address blocks using CIDR notation, separated by commas.

- **Subscribe to**: **Session Directory Topic**: Check this box to subscribe to user session information from the ISE server. Includes SGT and endpoint metadata. **SXP Topic**: Check this box to subscribe to SXP mappings from the ISE server.
- **Proxy**: You can optionally choose either a managed device or a proxy sequence to communicate with ISE/ISE-PIC if Security Cloud Control is unable to do so. For example, your Security Cloud Control might be in a public cloud but the ISE/ISE-PIC server might be on an internal intranet.



Note This version of the system does not support filtering using IPv6 addresses, regardless of your ISE version.

Troubleshoot Cisco ISE/ISE-PIC or Cisco TrustSec issues

This reference helps you troubleshoot issues with Cisco TrustSec, ISE, or ISE-PIC integration to ensure Security Group Tag propagation, user mapping, and certificate requirements are correctly configured for Cloud-Delivered Firewall Management Center and Secure Firewall Management Center deployments.

Troubleshoot Cisco TrustSec issues

A device interface can be configured to propagate Security Group Tags (SGTs) either from ISE/ISE-PIC or from a Cisco device on the network (referred to as Cisco TrustSec.) On the device management page (**Devices** > **Device Management**), the **Propagate Security Group Tag** check box for an interface is checked after a device reboot. If you do not want the interface to propagate TrustSec data, uncheck the box.

Troubleshoot ISE/ISE-PIC issues

For other related troubleshooting information, see Troubleshoot realms and user downloads, on page 2342 and Troubleshoot user control, on page 2526.

If you experience issues with the ISE or ISE-PIC connection, check these items:

- The pxGrid Identity Mapping feature in ISE must be enabled before you can successfully integrate ISE with the system.
- When the primary server fails, you must manually promote the secondary to primary; there is no automatic failover.
- Before a connection between the ISE server and the Cloud-Delivered Firewall Management Center succeeds, you must manually approve the clients in ISE. (Typically, there are two clients: one for the connection test and another for ISE agent.)

  You can also enable **Automatically approve new accounts** in ISE as discussed in the chapter on Managing users and external identity sources in the Identity Services Engine administrator guides.
- The **pxGrid Client Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.

- The time on your ISE server must be synchronized with the time on the Cloud-Delivered Firewall Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.
- If your deployment includes a primary and a secondary pxGrid node,
  - The certificates for both nodes must be signed by the same certificate authority.
  - The ports used by the host name must be reachable by both the ISE server and by the Cloud-Delivered Firewall Management Center.
- If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.

To exclude subnets from receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE, use the **configure identity-subnet-filter** {**add** | **remove**} command. You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.

If you experience issues with user data reported by ISE or ISE-PIC, note these items:

- After the system detects activity from an ISE user whose data is not yet in the database, the system retrieves information about them from the server. Activity seen by the ISE user is *not* handled by access control rules, and is *not* displayed in the web interface until the system successfully retrieves information about them in a user download.
- You cannot perform user control on ISE users who were authenticated by an LDAP, RADIUS, or RSA domain controller.
- The Cloud-Delivered Firewall Management Center does not receive user data for ISE Guest Services users.
- If ISE monitors the same users as TS Agent, the Cloud-Delivered Firewall Management Center prioritizes the TS Agent data. If the TS Agent and ISE report identical activity from the same IP address, only the TS Agent data is logged to the Cloud-Delivered Firewall Management Center.
- Your ISE version and configuration impact how you can use ISE in the system. For more information, see The ISE/ISE-PIC identity source, on page 2373.
- If you have Cloud-Delivered Firewall Management Center high availability configured and the primary fails, see the section on ISE and high availability in Best practices for ISE/ISE-PIC integration, on page 2376.
- ISE-PIC does not provide ISE attribute data.
- ISE-PIC cannot perform ISE ANC remediations.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, refer to RFC 959.

If you experience issues with supported functionality, see The ISE/ISE-PIC identity source, on page 2373 for more information about version compatibility.

ISE/ISE-PIC user timeout

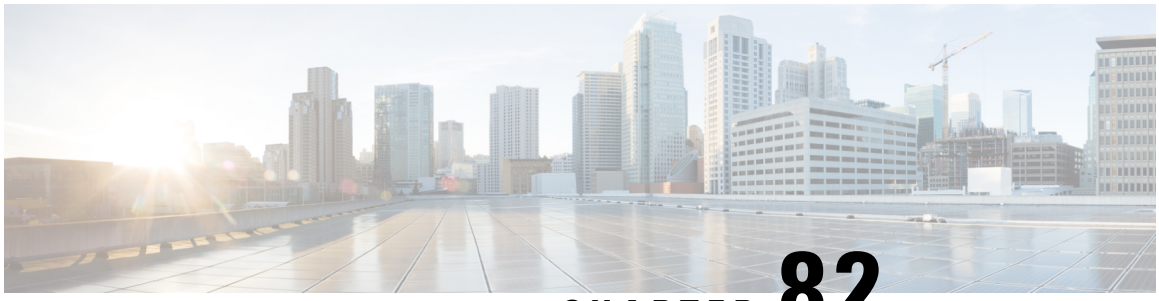
If you're setting up ISE/ISE-PIC without a realm, be aware there is a user session timeout that affects how users are seen by the Cloud-Delivered Firewall Management Center.

History for ISE/ISE-PIC

This reference provides the historical development of ISE and ISE-PIC integration features across different software versions.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Quick configuration	20241127	7.6.0	You can optionally configure ISE using only the user name and password of a user in the External RESTful Services (ERS) Operator group. (*This feature is available for ISE only, not ISE-PIC.*) **Upgrade impact**. Any ISE or ISE-PIC identity source you created before upgrade is still available on the **Advanced Configuration (Old)** tab page. Quick configuration affects only new ISE identity sources created after the upgrade. New/modified screens: **Integration** > **Other Integrations** > **Identity Sources** > **Identity Services Engine**. There are two tab pages: **Quick Configuration (New)** and **Advanced Configuration (Old)**. New/modified CLI commands: none
Proxy	Any	7.2.0	One or more managed devices that can communicate with Security Cloud Control in the event Security Cloud Control cannot communicate with the ISE/ISE-PIC server. **New/updated screen**: **Integration** > **Other Integrations** > **Realms** > **Proxy Sequence**
pxGrid 2.0 is the default for supported ISE/ISE-PIC versions	Any	6.7.0	Note: • Supported ISE/ISE-PIC versions: 2.6 patch 6 or later, 2.7 patch 2 or later • Adaptive Network Control (ANC) policies replace Endpoint Protection Service (EPS) remediations. If you have EPS policies configured in the Cloud-Delivered Firewall Management Center, you must migrate them to use ANC.
Optionally exclude subnets from receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE. You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.	Any	6.7.0	New command: **configure identity-subnet-filter** {**add** | **remove**}

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Destination Security Group Tag matching (SGT)	Any	6.5.0	Feature introduced. Enables you to use ISE SGT tags for both source and destination matching criteria in access control rules. SGT tags are tag-to-host/network mappings obtained by ISE. New/modified screens: • New options to configure Destination SGT matching: **System** > **Integration** > **Identity Sources** > **ISE/ISE-PIC** • **Session Directory Topic**: Subscribe to ISE user session information. • **SXP Topic**: Subscribe to SGT tag updates on the ISE server. • New and renamed columns in **Analysis** > **Connections** > **Events** • Renamed: Security Groups Tags renamed to Source SGT • New: Destination SGT
Integration with ISE-PIC	Any	6.2.1	You can now use data from ISE-PIC.
SGT tags for user control.	Any	6.2.0	You no longer need to create a realm or identity policy to perform user control based on ISE Security Group Tag (SGT) data.
Integration with ISE.	Any	6.0	Feature introduced. By subscribing to Cisco's Platform Exchange Grid (PxGrid), the Firepower Management Center can download additional user data, device type data, device location data, and Security Group Tags (SGTs) —a method used by ISE to provide network access control).



CHAPTER 82

Identity Source: Passive Identity Agent

The following topics discuss how to configure and use the passive identity agent.

- The passive identity agent identity source, on page 2403
- Deploy the passive identity agent, on page 2404
- How to configure a passive identity agent identity source, on page 2409
- Configure the passive identity agent, on page 2410
- Monitor the passive identity agent, on page 2426
- Manage the passive identity agent, on page 2427
- Troubleshoot the passive identity agent, on page 2429
- Security requirements for the passive identity agent, on page 2430
- Internet access requirements for the passive identity agent, on page 2431
- History for the passive identity agent, on page 2431

The passive identity agent identity source

The passive identity agent identity source sends session data from Microsoft Active Directory (AD) to the Cloud-Delivered Firewall Management Center. All you need is a supported Microsoft AD setup as described in Realms and realm sequences, on page 2273.

Passive identity agent roles

The passive identity agent supports these roles: Standalone, Primary, and Secondary.

Note You do not need to configure the Cisco Identity Services Engine (ISE) to use this identity source.

For more information on these roles, refer to About passive identity agent roles.

Passive identity agent system requirements

The passive identity agent supports Windows Server 2016, 2019, 2022, and 2025 on Active Directory servers, and Windows 11 or later on domain-joined Windows clients. The system time must be synchronized across the Cloud-Delivered Firewall Management Center, domain controllers, and the agent host.

Refer to Passive identity agent system requirements for more information.

Passive identity agent limitations

The passive identity agent has these limitations:

- Up to 10 agents simultaneously
- A single passive identity agent identity source monitors up to 50 AD directories
- Up to 300,000 concurrent user sessions
- IPv6 addresses are *not* supported (passive identity agent 1.0)
- IPv6 addresses are supported (passive identity agent 1.1)
- Cloud-Delivered Firewall Management Center receives user login information from the Passive Identity Agent but does not track or process user logoff events. As a result, Cloud-Delivered Firewall Management Center does not detect user session termination and continues to show user information until either the realm times out or the administrator deletes the user session.

Deploy the passive identity agent

For information about deployment options, see Deploy the passive identity agent, on page 2404.

Note We recommend you use the latest version of the passive identity agent.

To see the available versions, go to software.cisco.com.

To upgrade the passive identity agent, see Upgrade the passive identity agent software, on page 2426.

Deploy the passive identity agent

Deploy the passive identity agent software on any machine that is part of a Microsoft Active Directory (AD) domain that you want to use for user awareness and control.

Installation locations and agent types

You can install the passive identity agent software on any of the following machines:

- The Microsoft Active Directory server
- A domain controller
- A client connected to the network that is neither the directory server nor a domain controller

Any particular passive identity agent can monitor one or several Active Directory domain controllers in the same domain.

The machine on which the passive identity agent is installed must communicate with the Cloud-Delivered Firewall Management Center using the TLS/SSL protocol. For more information, see Internet access requirements for the passive identity agent, on page 2431.

You can configure the following types of agents on the Microsoft AD directory server, domain controller, or on any client connected to the domain:

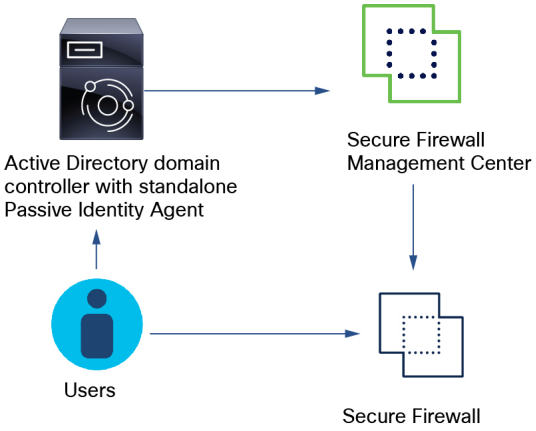
- Standalone agent: One agent that can monitor one or several Active Directory domain controllers in the same domain.
- Primary agent and secondary agent: Both can monitor one or several AD domain controllers in the same domain. To provide redundancy, you can install a primary agent and a secondary agent on different machines. The primary agent is responsible for communicating with the Cloud-Delivered Firewall Management Center. If communication fails, the secondary agent takes over.

Related Topics

Simple passive identity agent deployment, on page 2405
Single passive identity agent monitoring multiple domain controllers, on page 2405
Multiple passive identity agents monitoring multiple domain controllers, on page 2406
Configure Passive identity agent primary and secondary agent deployments, on page 2408

Simple passive identity agent deployment

The following diagram shows the simplest passive identity agent deployment.



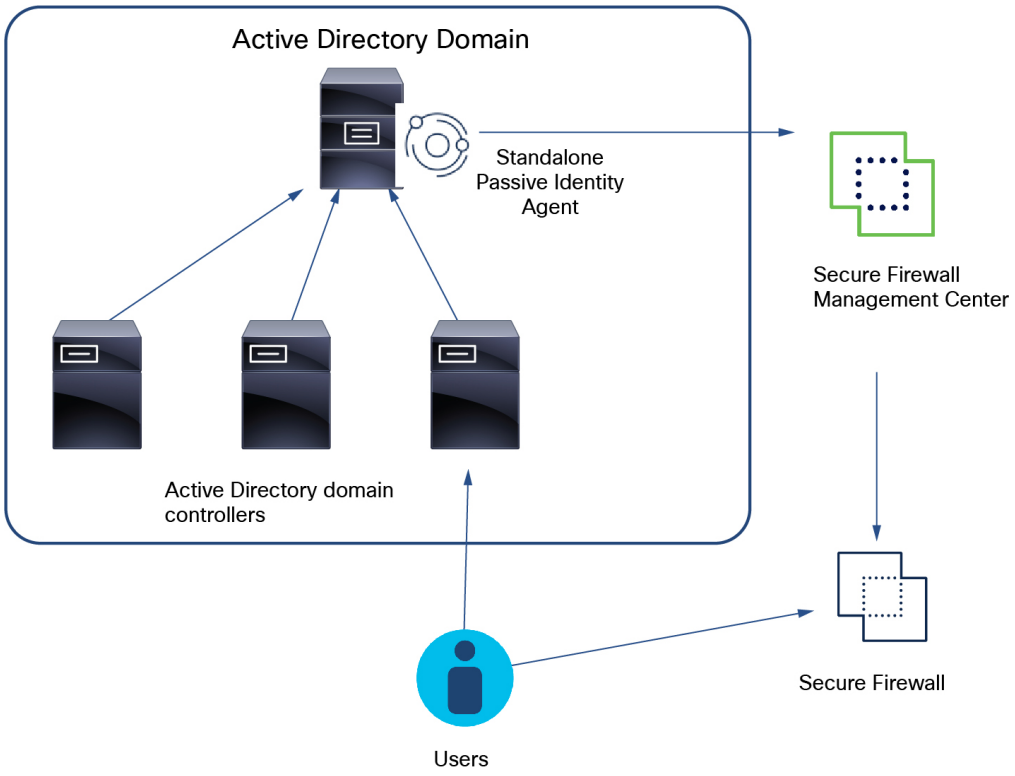
Deployment architecture

In this example, a standalone passive identity agent is installed on the AD domain controller. Users log in and out of the AD domain. The agent sends user name and IP address information to the Cloud-Delivered Firewall Management Center. As users access the network, access control and identity policies deployed to the Secure Firewall Threat Defense determine whether or not access is allowed and how access is granted.

You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor.

Single passive identity agent monitoring multiple domain controllers

The diagram shows a standalone passive identity agent that monitors several AD domain controllers.



Architecture details

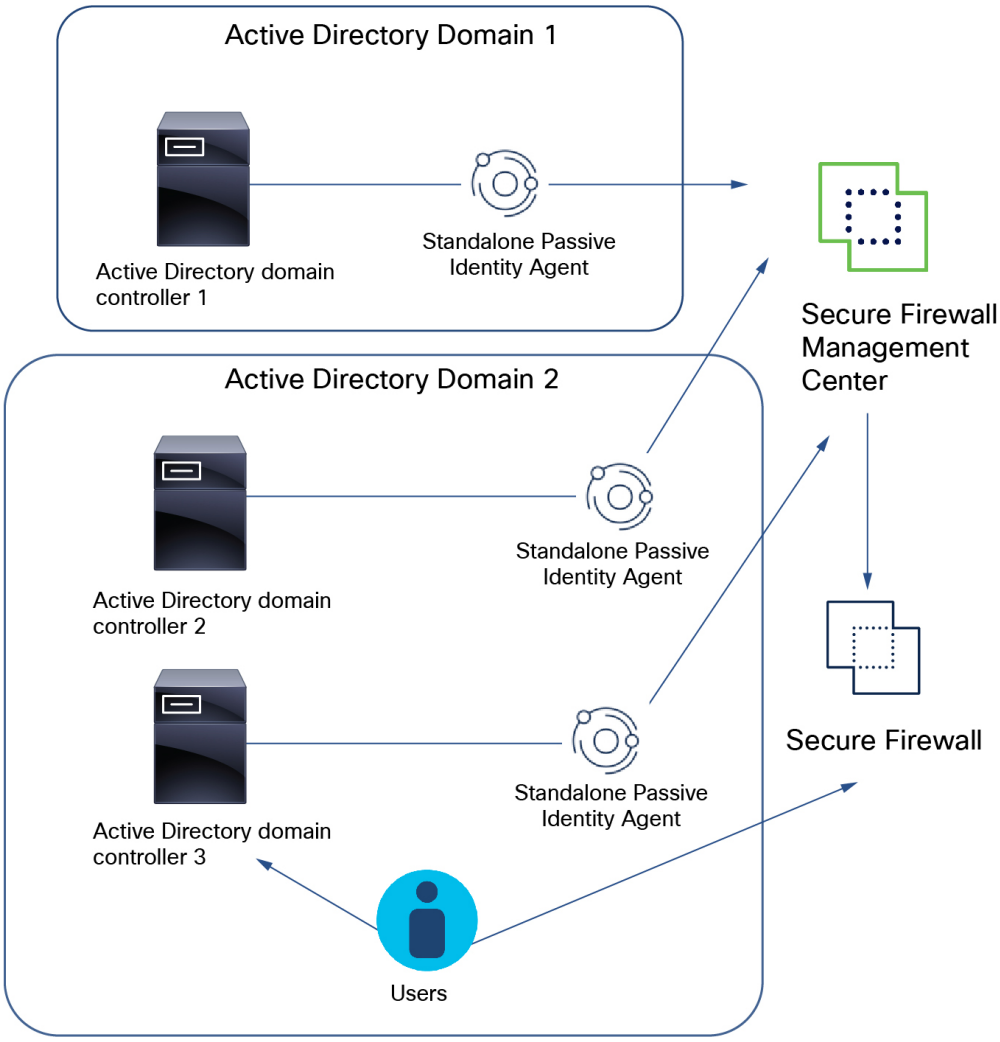
In this diagram, the standalone passive identity agent is installed on a client attached to the AD domain (or on the domain controller itself). Users log in to any domain controller and the agent sends user and IP address information to the Cloud-Delivered Firewall Management Center. As users access the network, access control and identity policies deployed to the Secure Firewall Threat Defense determine whether access is allowed and how access is permitted.

You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor.

Multiple passive identity agents monitoring multiple domain controllers

This figure illustrates how standalone agents monitor multiple AD domain controllers.

- In AD domain 1, a standalone passive identity agent installed on a machine attached to AD domain controller 1 sends user and IP address mapping data to the .
- In AD domain 2, standalone agents installed on domain controllers 1 and 2 send user and IP address mapping data to the .



You can install a passive identity agent on an AD domain controller, a directory server, or on any client connected to the domain you want to monitor.

Deployment steps

The preceding figure shows three passive identity agents, each configured as a standalone. To do this:

1. Create two Microsoft AD realms: one for each AD domain.

   See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

2. For AD domain 2, create two directories, one for each domain controller.
3. Install the Passive Identity Agent software on a client that can log in to the domain.

   Configure each passive identity agent individually to communicate with the on which you configure the passive identity agent source.

   See Install the passive identity agent software, on page 2422.

4. Create the passive identity agent identity source.

See Create a primary or secondary passive identity agent identity source, on page 2414.

Configure Passive identity agent primary and secondary agent deployments

Configure primary and secondary passive identity agents to provide redundancy and avoid a single point of failure in user identity monitoring.

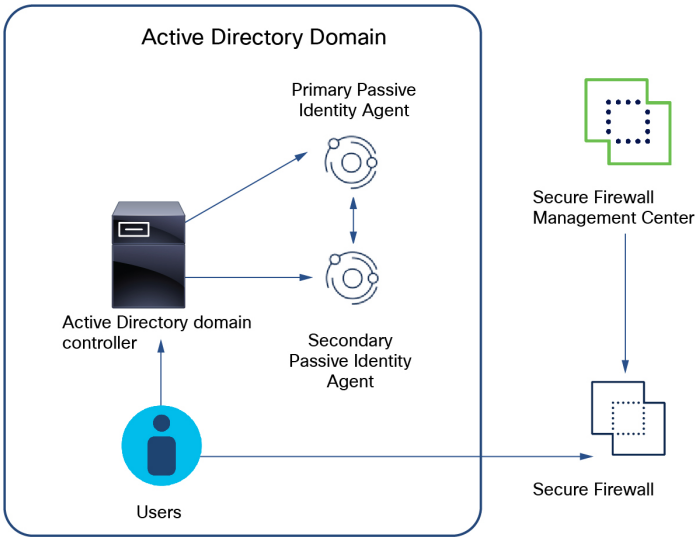
To provide redundancy and to avoid a single point of failure, you can configure primary and secondary passive identity agents in any of the ways shown in this topic.

You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor.

Procedure

Step 1 For a single AD domain controller with primary and secondary agents deployment:

The following figure shows how to set up primary and secondary passive identity agents on one AD domain controller. If the primary agent fails, the secondary takes over.



a) Create a Microsoft AD realm that has one directory for the domain controller.

See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

b) Install the passive identity agent software on any two network machines connected to the domain controller.

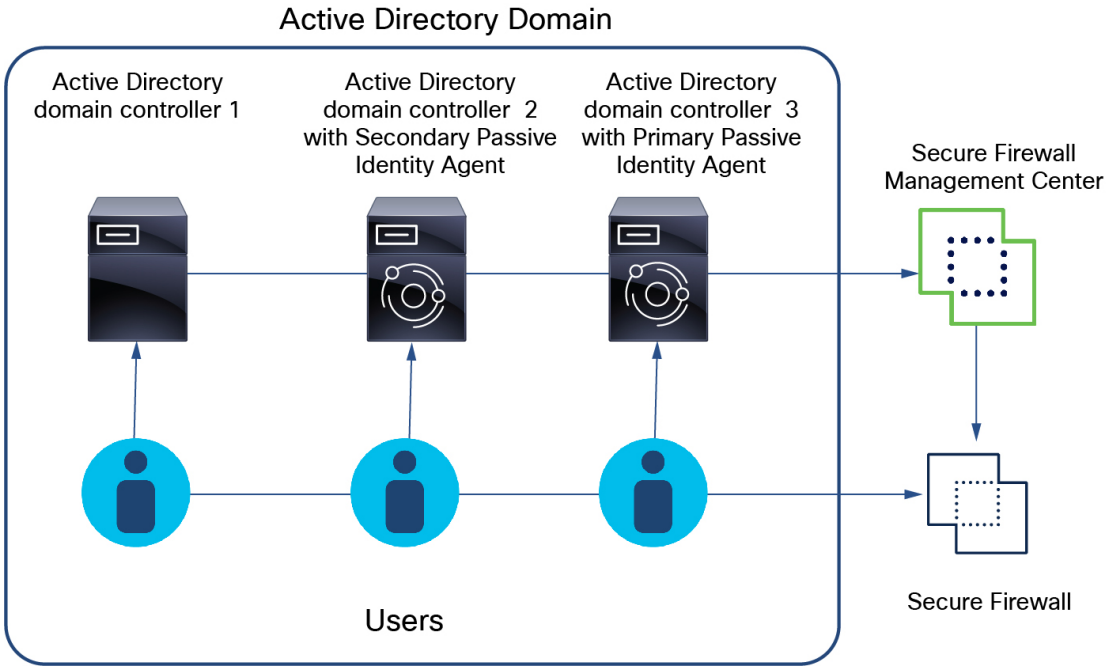
Configure each passive identity agent individually to communicate with the Cloud-Delivered Firewall Management Center on which you configure the passive identity agent source.

See Install the passive identity agent software, on page 2422.

c) Create the identity source.

See Create a primary or secondary passive identity agent identity source, on page 2414.

Step 2 For multiple AD domain controllers with primary and secondary agents deployment.



The preceding figure shows how to configure primary and secondary agents to monitor three AD domain controllers. If the primary agent fails, the secondary agent takes over.

a) Create a Microsoft AD realm that has one directory for the domain controller.

See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

b) Install the passive identity agent software on any machine connected to the domain controller.

Configure each passive identity agent individually to communicate with the Cloud-Delivered Firewall Management Center on which you configure the passive identity agent source.

See Install the passive identity agent software, on page 2422.

c) Create the identity source.

See Create a primary or secondary passive identity agent identity source, on page 2414.

How to configure a passive identity agent identity source

Configure the passive identity agent identity source to enable user identity monitoring and enhance access control in your network environment.

This provides high-level tasks required to configure the passive identity agent identity source in the Cloud-Delivered Firewall Management Center and to deploy agent software to your Microsoft Active Directory (AD) servers.

Before you begin

Follow these steps to create a passive identity agent identity source:

Procedure

Step 1 (Optional.) Enable the Dynamic Attributes Connector.

Step 2 Create a realm for your Microsoft AD domain and domain controllers.

Realms are connections between the Cloud-Delivered Firewall Management Center and the user accounts on the servers you monitor. They specify the connection and authentication filter settings for the server.

For more information, see Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

Step 3 Create a passive identity agent identity source.

The identity source allows the Cloud-Delivered Firewall Management Center and passive identity agent to communicate with each other. Create standalone, primary, or secondary agents, to match your requirements.

For more information, refer to:

- Passive identity agent roles, on page 2417
- Create a passive identity agent identity source, on page 2411

Step 4 Get the API token for the passive identity agent to authenticate with Cloud-Delivered Firewall Management Center.

Refer to Get an API token for the passive identity agent, on page 2418.

Step 5 Install the passive identity agent software.

The installation method for the agent varies according to your deployment.

You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor.

For more information, see:

- Deploy the passive identity agent, on page 2404
- Install the passive identity agent software, on page 2422

What to do next

Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

Configure the passive identity agent

These topics discuss how to configure the passive identity agent.

Related Topics

Create an LDAP realm or an Active Directory realm and realm directory, on page 2325

Create a passive identity agent identity source, on page 2411
Install the passive identity agent software, on page 2422

Enable the dynamic attributes connector

Enable the dynamic attributes connector to allow objects from cloud networking products to be used in Cloud-Delivered Firewall Management Center access control and DNS rules.

This task discusses how to enable the dynamic attributes connector in the Cloud-Delivered Firewall Management Center. The dynamic attributes connector is an integration that enables objects from cloud networking products to be used in Cloud-Delivered Firewall Management Center access control and DNS rules.

Procedure

Step 1 Log in to the Cloud-Delivered Firewall Management Center if you have not done so already.

Step 2 Click **Administration** > **Dynamic Attributes Connector**.

Step 3 Slide to **Enabled**.

Step 4 The system displays messages when the dynamic attributes connector is enabled.

In the event of errors, try again. If errors persist, contact Cisco TAC.

See Connectors, on page 2631.

Create a Microsoft Active Directory realm

passive identity agent requires you to create a Microsoft Active Directory (AD) realm and directories in the Cloud-Delivered Firewall Management Center.

Additional information

For detailed instructions on creating a Microsoft Active Directory realm and directories, see Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

Create a passive identity agent identity source

This task creates a passive identity agent that sends user session activity to the Cloud-Delivered Firewall Management Center.

Use this task to configure agent details, assign roles, and integrate with Microsoft AD for identity management and monitoring.

Before you begin

Complete the following:

- Review passive identity agent roles as discussed in Passive identity agent roles, on page 2417.
- Create a Microsoft AD realm as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

Follow these steps to create a passive identity agent identity source:

Procedure

Step 1 Log in to the Security Cloud Control.

Step 2 Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 3 Click **Passive Identity Agent**.

Step 4 If the dynamic attributes connector has not been enabled yet, you are prompted to do so.

For more information about enabling the dynamic attributes connector, see Enable the dynamic attributes connector, on page 2411.

Step 5 Click **Create Agent**.

Step 6 In the Configure Agent dialog box, enter the following information:

Item	Description
Name	Enter a unique name to identify this passive identity agent.
Description	Enter an optional description.
Role	Click one of the following: • **Primary**: The agent responsible for communicating with the Cloud-Delivered Firewall Management Center. Not available if you choose **Standalone**. • **Secondary**: Becomes the primary if the primary loses contact with the Cloud-Delivered Firewall Management Center. Not available if you choose **Standalone**. • **Standalone**: If there is only one passive identity agent. For more information about roles, see Passive identity agent roles, on page 2417.

Step 7 Continue with:

- Create a standalone passive identity agent identity source, on page 2412
- Create a primary or secondary passive identity agent identity source, on page 2414

Create a standalone passive identity agent identity source

Configure a standalone passive identity agent to enable identity management and user control.

This task discusses how to configure a standalone passive identity agent.

Before you begin

Complete the tasks discussed in Create a passive identity agent identity source, on page 2411.

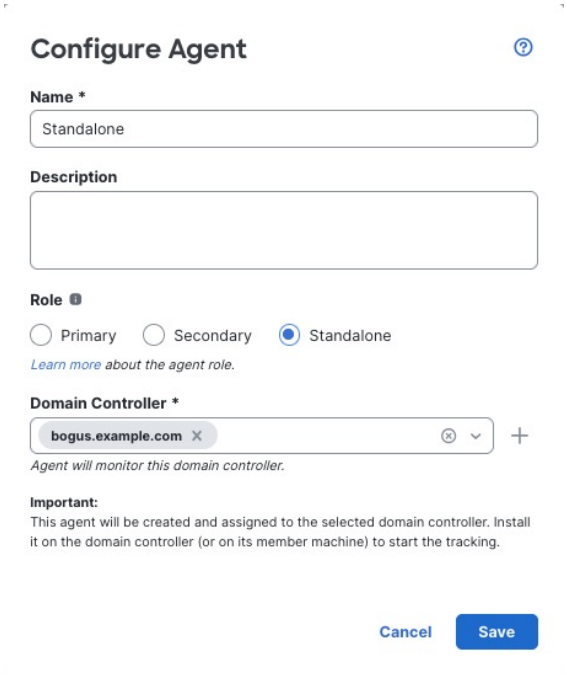
Follow these steps to create a standalone passive identity agent identity source:

Procedure

Step 1 In the Configure Agent dialog box, enter the following information:

Item	Description
Role	Click **Standalone**.
Domain Controller	From the list, select the check box next to each domain controller that has a passive identity agent you wish to use for identity management and user control. (Optional.) Click **Add** (+) to add a new one.

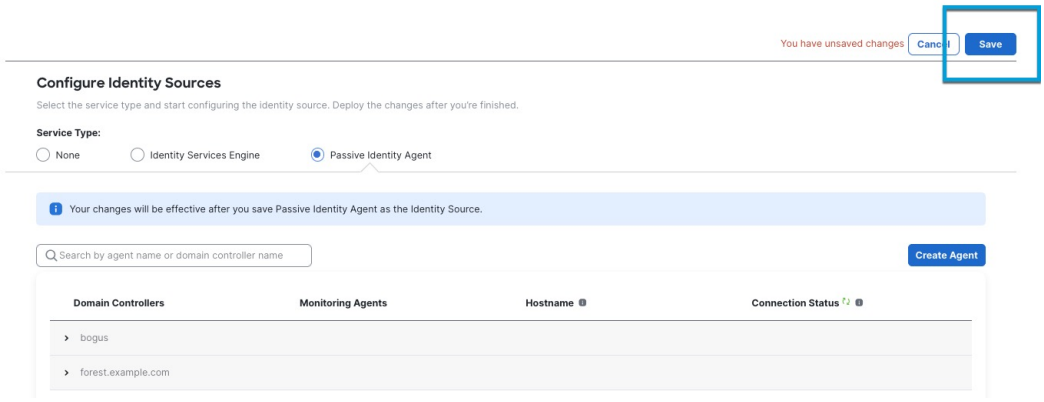
This figure shows an example of a standalone passive identity agent identity source.



Step 2 In the Configure Agent dialog box, click **Save**.

Step 3 In the top right corner of the page, click **Save**.

This figure shows an example.



Note
The passive identity agent won't be active until you create a user and install the software.

What to do next

- See Get an API token for the passive identity agent, on page 2418
- See Passive identity agent installation, on page 2418

Create a primary or secondary passive identity agent identity source

This task enables you to configure a primary and secondary passive identity agent to manage user identities and ensure continuous communication with the firewall. You will assign agent roles, specify host details, and enable domain controller monitoring for identity management.

The following task continues from Create a passive identity agent identity source, on page 2411.

Before you begin

Complete the tasks discussed in Create a passive identity agent identity source, on page 2411.

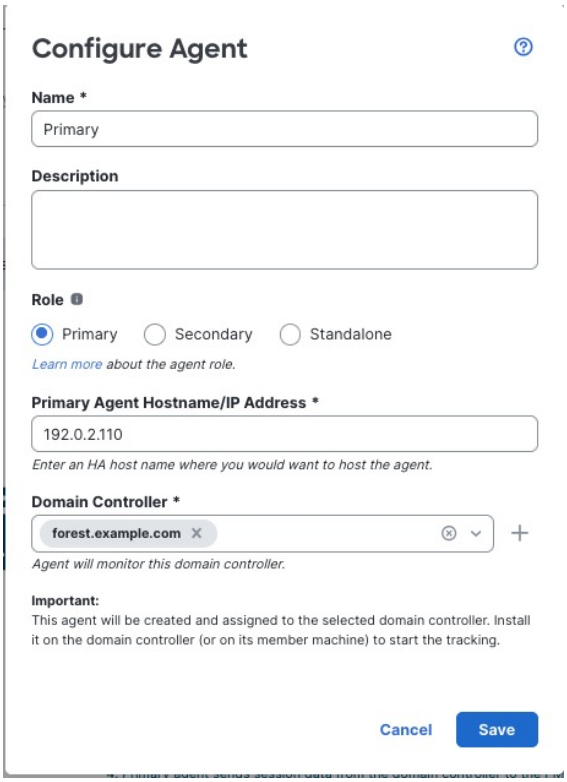
Follow these steps to create a primary or secondary passive identity agent identity source:

Procedure

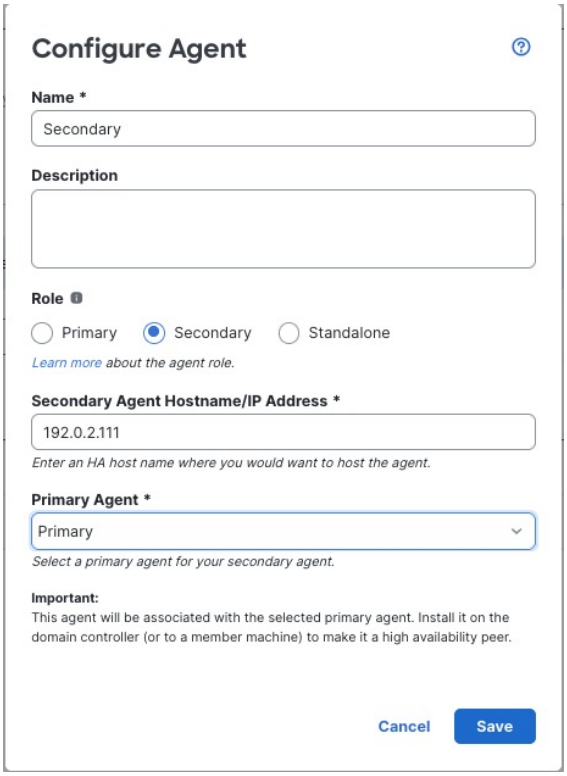
Step 1 In the Configure Agent dialog box, enter the following information:

Item	Description
Role	Click one of the following: • **Primary**: The agent responsible for communicating with the Cloud-Delivered Firewall Management Center. • **Secondary**: Becomes the primary if the primary loses contact with the Cloud-Delivered Firewall Management Center. For more information about roles, see Passive identity agent roles, on page 2417.
Primary Agent Hostname/IP Address	(Primary agent only.) Enter the fully qualified domain name or IP address of the server on which the primary passive identity agent is installed. The passive identity agent version 1.0 supports IPv4 addresses and fully qualified domain names only. Version 1.1 supports IPv4, IPv6, and fully qualified domain names.
Secondary Agent Hostname/IP Address	(Secondary agent only.) Enter the fully qualified host name or IP address of the server on which the secondary passive identity agent is installed. The passive identity agent version 1.0 supports IPv4 addresses and fully qualified domain names only. Version 1.1 supports IPv4, IPv6, and fully qualified domain names.
Primary Agent	(Secondary agent only.) From the list, click the name of the primary passive identity agent.
Domain Controller	(Primary agent only.) From the list, select the check box next to each domain controller that has a passive identity agent you wish to use for identity management and user control.

The following figure shows an example of a primary agent:



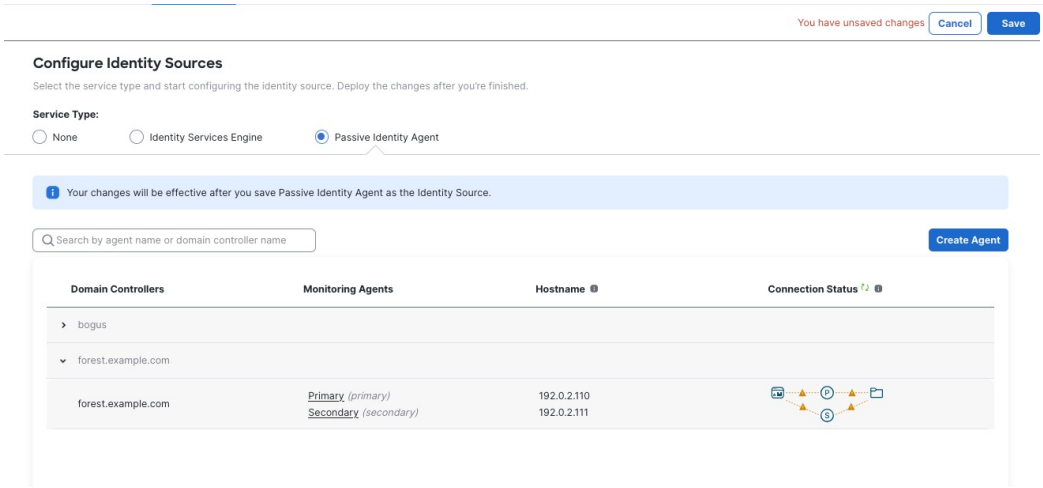
The following figure shows an example of a secondary agent:



Step 2 In the Configure Agent dialog box, click **Save**.

Step 3 In the top right corner of the page, click **Save**.

The following figure shows an example.



Note
The passive identity agent won't be active until you create a user and install the software.

What to do next

- See Get an API token for the passive identity agent, on page 2418
- See Passive identity agent installation, on page 2418

Passive identity agent roles

A passive identity agent role determines how the agent functions in Active Directory environments for user and group monitoring.

Role types

The passive identity agent has the following roles:

- Standalone: A passive identity agent that is not part of a redundant pair. A standalone agent can read users and groups from multiple Active Directory servers and domain controllers, provided the software is installed on all of them.
- Primary: (Primary agent in a redundant pair.) Can be installed on a Microsoft AD domain controller, directory server, or any network client.

  Handles all communication with the Cloud-Delivered Firewall Management Center unless it stops communicating, in which case communication is handled by secondary agents.
- Secondary: (Secondary, or backup, agent in a redundant pair.) Can be installed on a Microsoft AD domain controller, directory server, or any network client.

Monitors the health of the primary agent and takes over if the primary agent stops communicating with the Cloud-Delivered Firewall Management Center.

The passive identity agent can monitor several AD domain controllers that are part of the same domain.

Get an API token for the passive identity agent

Use this procedure to obtain the API token that enables the passive identity agent to authenticate with the Cloud-Delivered Firewall Management Center.

This task applies *only* to using the passive identity agent with the Cloud-Delivered Firewall Management Center.

Required role:

- Super Admin
- Admin

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Administration** > **Firewall Management Center**.

Step 3 Check the box next to the Cloud-Delivered Firewall Management Center with which you want to use the passive identity agent.

Step 4 In the right pane, click **Passive Identity**.

Step 5 Copy the token to the clipboard.

What to do next

Install the passive identity agent software, on page 2422.

Passive identity agent installation

The passive identity agent installation process requires specific prerequisites and tasks to deploy the agent in your environment.

Installation requirements

These topics discuss prerequisites and tasks required to install the passive identity agent.

Note We recommend you use the latest version of the passive identity agent. To see the available versions, go to software.cisco.com. To upgrade the passive identity agent, see Upgrade the passive identity agent software, on page 2426

Prerequisites to install the passive identity agent

You must complete all of the tasks before you install the passive identity agent software.

Passive identity agent system requirements

This reference provides the system and software requirements for installing the passive identity agent, including supported Windows versions, synchronized system clocks, and necessary firewall software versions.

- If you install the passive identity agent on a Windows Active Directory server, the server must run one of the following versions:
  - Microsoft Windows Server 2016
  - Microsoft Windows Server 2019
  - Microsoft Windows Server 2022
- If you install the passive identity agent on a Windows client attached to the domain, the client must run Windows 11 or later.
- The system clock on all systems must be synchronized. We strongly recommend using the same NTP servers on all of them. This means:
  - The Security Cloud Control.

    For more information, see Configure NTP Server.
  - All Windows Active Directory servers and domain controllers.
  - The machine on which the passive identity agent is installed.
- Security Cloud Control must run November 8, 2024 or later.
- You must enable Snort 3 on the Secure Firewall Threat Defense devices.

Enable the windows event viewer to log kerberos authentication attempts

This task enables the Windows Event Viewer to log successful and unsuccessful Kerberos authentication attempts so that the passive identity agent can function properly by reading user sessions from the Event Viewer.

The passive identity agent reads user sessions from the Event Viewer so this setting is required for the passive identity agent to function properly.

For more information, see System audit policy recommendations on learn.microsoft.com.

Procedure

Step 1 Log in to the Active Directory Server as an administrator and open a DOS command prompt.

Step 2 Enter `gpmc.msc` to start the Group Policy Management Editor.

Step 3 If necessary, create a new Group Policy Object (GPO); if one already exists, edit it.

For more information about creating a GPO, see a resource like Create a Group Policy Object on learn.microsoft.com.

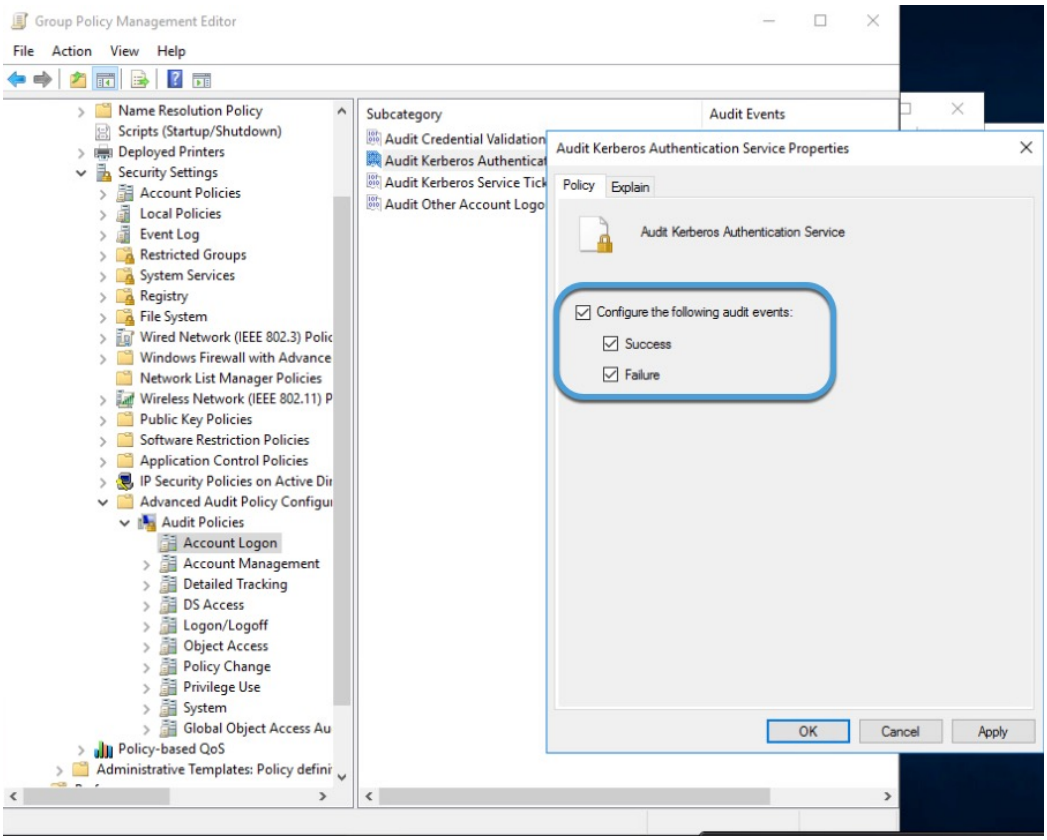
Step 4 In your GPO, expand **Computer Configuration** > **Policies** > **Windows Settings** > **Security Settings** > **Advanced Policy Configuration** > **Audit Policies**.

Step 5 Click **Account Logon**.

Step 6 In the right pane, double-click **Audit Kerberos Authentication Service**.

Step 7 In the dialog box that is displayed, select all checkboxes which enables the system to log successes and failures.

The following figure shows an example.



Step 8 Follow the prompts on your screen to save the changes.

Step 9 (Optional.) To update GPO immediately, enter **`gpupdate /force`** in your DOS command prompt window.

What to do next

See Add the Active Directory user to groups, on page 2420.

Add the Active Directory user to groups

Use this procedure to give the Active Directory and passive identity agent service user sufficient privileges to Active Directory.

To function normally, the passive identity agent must be able to connect to the domain and to read the Windows Event Log. This topic discusses how to give the proper privileges to:

- The passive identity agent service user.

- Active Directory user (namely, the **Directory Username** user in the Active Directory realm on the Cloud-Delivered Firewall Management Center).

Before you begin

You must be a Microsoft Server administrator familiar with how to add a user to a group and how to set a Windows service to run as a specific user.

Procedure

Step 1 Log in as an administrator to the system on which the passive identity agent is running.

You can log in to any of the following:

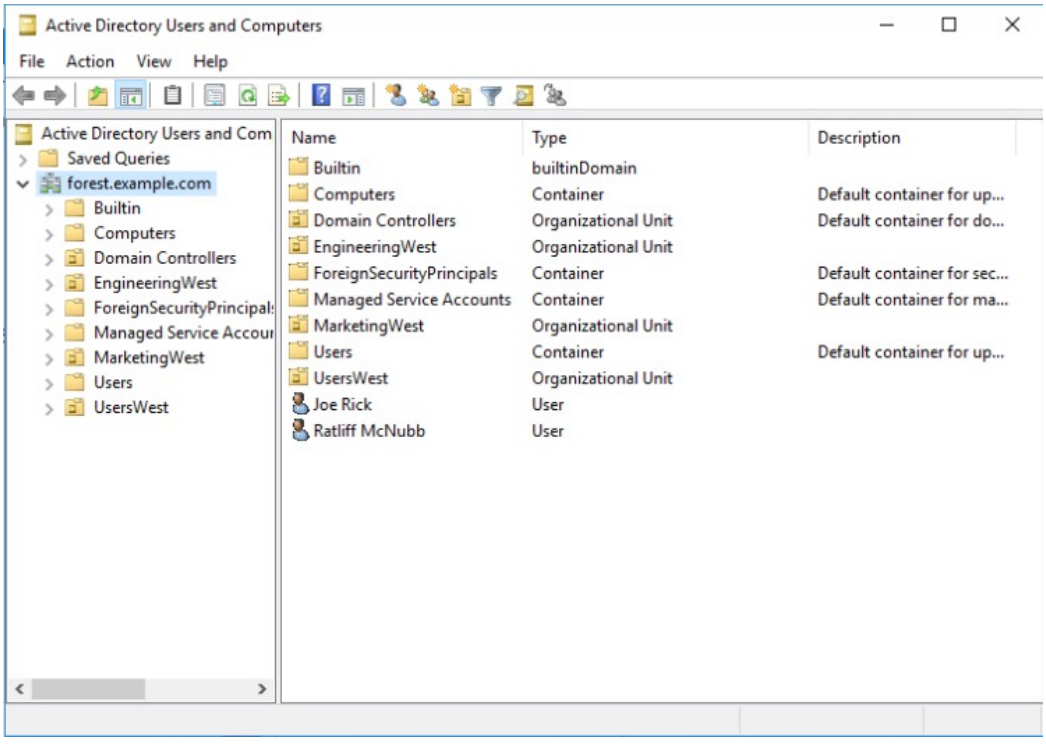
- The domain controller.
- The Active Directory server.

Step 2 Start the Server Manager.

Step 3 Click **Tools** > **Active Directory Users and Computers**.

Step 4 Under Active Directory Users and Computers, expand the forest in which the directory user is defined.

The following figure shows an example.

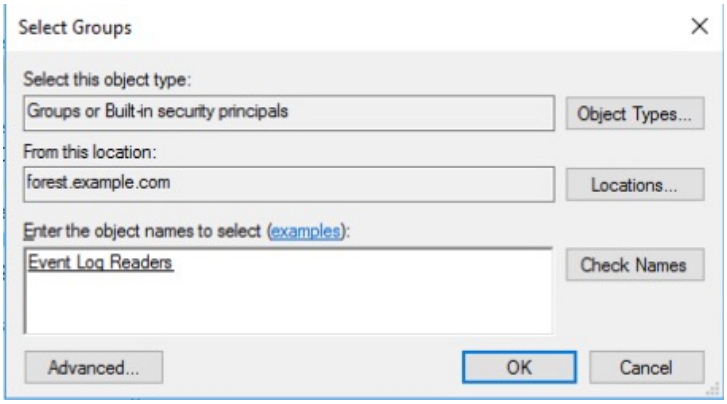


Step 5 Expand the organization unit or group to reveal the directory user. (You can create a new user by clicking **New** > **User**).

Step 6 Right-click the directory user and click **Add to a group**.

Step 7 In the Select Groups dialog box, enter **`Event Log Readers`** and click **Check Names**.

The following figure shows an example.



Step 8 Repeat the preceding tasks to add the user to the Domain Users group.

Step 9 In the Add Groups dialog box, click **OK**.

What to do next

See Install the passive identity agent software, on page 2422.

Install the passive identity agent software

Install the passive identity agent software to enable user identity monitoring and integration with your firewall management center.

This task discusses how to install the passive identity agent software. For a simple installation, you can install it on your Microsoft Active Directory (AD) domain controller; for other options, see Deploy the passive identity agent, on page 2404.

Note We recommend you use the latest version of the passive identity agent. To see the available versions, go to software.cisco.com. To upgrade the passive identity agent, see Upgrade the passive identity agent software, on page 2426

Before you begin

See Get an API token for the passive identity agent, on page 2418.

Complete all of these tasks:

- Enable the windows event viewer to log kerberos authentication attempts, on page 2419
- Add the Active Directory user to groups, on page 2420
- Add Log On to the Passive Identity Agent Service, on page 2424

Follow these steps to install the Passive Identity Agent Software:

Procedure

Step 1 Download the passive identity agent from software.cisco.com.

Step 2 Log in as a member of the Administrators group to the machine on which to install the passive identity agent.

Step 3 Double-click **CiscoPassiveIdentityAgentInstaller-1.1.1.msi** and click **Next**.

Step 4 Choose a folder in which to install the passive identity agent and click **Next**.

The default installation folder is **Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent**.

Step 5 Click **Install**.

Step 6 When the installation is done, click **Finish** and optionally check the box to start the passive identity agent.

Step 7 When the passive identity agent starts, click the **On-Prem** tab if you are using the agent with an on-premises Cloud-Delivered Firewall Management Center (physical or virtual) or click the **Cloud** tab if you are using the agent with Cloud-Delivered Firewall Management Center.

Step 8 In the Cisco Passive Agent dialog box, enter the following information:

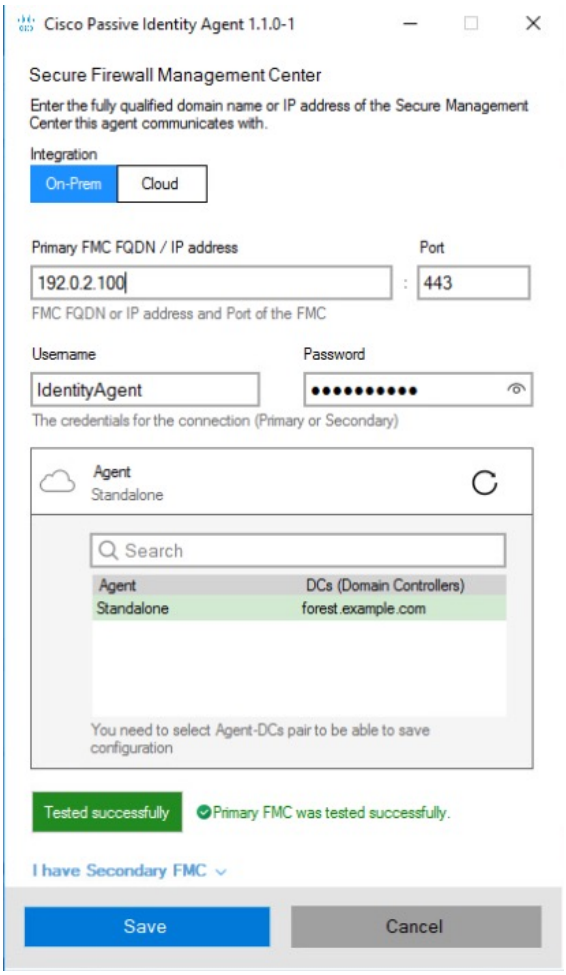
Item	Description
FMC FQDN / IP Address	Enter the fully qualified domain name, IPv4, or IPv6 address of the Cloud-Delivered Firewall Management Center on which you created the passive identity agent identity source.
Token	Enter the API token you found in Get an API token for the passive identity agent, on page 2418.
Agent	Click the list to locate the domain controller of the passive identity agent you created previously on the Cloud-Delivered Firewall Management Center.

Step 9 Click the **Agent** list.

a) From the list, click the name of the domain controller to monitor.

b) Click **Test**.

The following figure shows an example.



c) Only if the test succeeds, click **Save**.

What to do next

See Add Log On to the Passive Identity Agent Service, on page 2424.

Add Log On to the Passive Identity Agent Service

Use this procedure to enable the passive identity agent service to run as the Active Directory user. This user is the **Directory Username** user in the Active Directory realm on the Cloud-Delivered Firewall Management Center.

This task is optional but recommended. It ensures that the passive identity agent service runs with the minimal permissions required to send login information to the Cloud-Delivered Firewall Management Center.

Before you begin

Complete the tasks in Add the Active Directory user to groups, on page 2420.

You must be a Microsoft Server administrator who knows how to add a user to a group and set a Windows service to run as a specific user.

Procedure

Step 1 Log in as an administrator to the system on which the passive identity agent is running.

You can log in to any of these:

- the domain controller
- the Active Directory server

Step 2 In the Windows search bar, enter **`Services`**.

Step 3 In the Services window, right-click **Cisco Passive Identity Agent**.

Step 4 Click **Properties**.

Step 5 In the Properties dialog box, click the **Log On** tab.

Step 6 Click **This account**.

Step 7 Click **Browse**. Follow the prompts on your screen to select the directory user.

Step 8 Enter the user's password in the provided fields.

Step 9 Click **Apply**.

What to do next

- Specify users to control and other options using an identity policy as described in Create an identity policy, on page 2507.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Deploy your identity and access control policies to managed devices as discussed in Deploy configuration changes, on page 202.
- Monitor user activity .

Uninstall the passive identity agent software

Remove the passive identity agent software from Microsoft AD servers when it is no longer needed or requires replacement.

This task discusses how to uninstall the passive identity agent software from your Microsoft AD servers.

Procedure

Step 1 Log in as an administrator to the machine on which the passive identity agent is installed.

Step 2 Search for Add or Remove Programs.

Step 3 Click **Cisco Passive Identity Agent**.

Step 4 Click **Uninstall**.

Step 5 You are required to confirm the uninstallation.

Upgrade the passive identity agent software

To upgrade to a newer version of the passive identity agent, first uninstall the earlier version. Then install the later version.

Refer to this information:

- Uninstall the passive identity agent software, on page 2425
- Passive identity agent installation, on page 2418

Monitor the passive identity agent

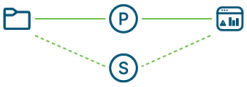
The passive identity agent indicates whether or not it can communicate with the Security Cloud Control and other agents if it's configured as primary-secondary. View the status at **Integration** > **Other Integrations** > **Identity Sources**.

Deployments

A standalone passive identity agent is represented as follows.



A primary-secondary pair is represented as follows.



This table explains the meaning of the indicators.

Object	Meaning
	Cloud-Delivered Firewall Management Center
	Standalone Passive Identity Agent
	Active Directory domain controller
	Primary agent
	Secondary agent

Status indicators and colors

The passive identity agent indicates status using lines (that indicate whether communication with the Cloud-Delivered Firewall Management Center is active or standby) and colors (that indicate whether or not communication is successful).

This table explains the meanings of lines and colors:

Object	Meaning
Solid line	The agent that is responsible for communicating with the Cloud-Delivered Firewall Management Center.
Dashed line	Primary/secondary configuration only. The agent that is acting as the backup agent. In the event of a communication failure between the active (solid line) agent, this agent communicates with the Cloud-Delivered Firewall Management Center.
Green	Agent communication is normal.
Amber 	Agent has never successfully communicated with the Cloud-Delivered Firewall Management Center. A newly created agent line is amber and remains so until configuration is complete.
Red 	Communication is failing. To resolve the issues: • Verify the network connections between agents and the Cloud-Delivered Firewall Management Center. • Complete the configuration for the system (Microsoft AD server, domain controllers, and the Cloud-Delivered Firewall Management Center). For more information, see How to configure a passive identity agent identity source, on page 2409.

Manage the passive identity agent

passive identity agent management includes editing previously configured agents and deleting existing agent configurations.

Platform-specific management options

Related Topics

Edit passive identity agents, on page 2428
Delete a standalone passive identity agent, on page 2428

Delete primary and secondary passive identity agents, on page 2428
Uninstall the passive identity agent software, on page 2425

Edit passive identity agents

Edit passive identity agents to update configuration settings and maintain accurate identity source integration for your network security environment.

This task discusses how to edit passive identity agents you previously configured in the Security Cloud Control.

Procedure

Step 1 Log in to the Security Cloud Control.

Step 2 Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 3 Click **Passive Identity Agent**.

Step 4 Click **Edit** (✎) next to the agent to edit.

Step 5 Make the desired changes.

Step 6 Click **Save**.

Delete a standalone passive identity agent

Delete a standalone passive identity agent to remove its association from the Secure Firewall Management Center and Security Cloud Control for enhanced identity source management.

This task discusses how to delete a standalone passive identity agent.

Procedure

Step 1 Log in to the Security Cloud Control.

Step 2 Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 3 Click **Passive Identity Agent**.

Step 4 Click **Edit** (✎) next to the agent to delete.

Step 5 Click **Delete**.

Step 6 Confirm the action when prompted.

Delete primary and secondary passive identity agents

This task allows you to remove unwanted or unused passive identity agents from your system configuration.

You must delete a secondary agent before you can delete a primary agent.

Procedure

Step 1 Log in to the Security Cloud Control.

Step 2 Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 3 Click **Passive Identity Agent**.

Step 4 Click **Edit** ( ) next to a secondary agent to delete.

Step 5 Click **Delete**.

Step 6 Confirm the action when prompted.

Step 7 If you wish to delete a primary agent, first delete all secondary agents.

Troubleshoot the passive identity agent

Troubleshoot the passive identity agent software on your Windows AD domain controller or directory server.

This task helps you investigate and resolve user session issues with the passive identity agent by utilizing various troubleshooting methods including log analysis and event viewer investigation.

Procedure

Step 1 Set the log level by opening **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent\CiscoPassiveIdentityAgentService.exe.config** in a text editor, save the file, and restart the Cisco Passive Identity Agent service.

By default, the passive identity agent logs at the INFO level.

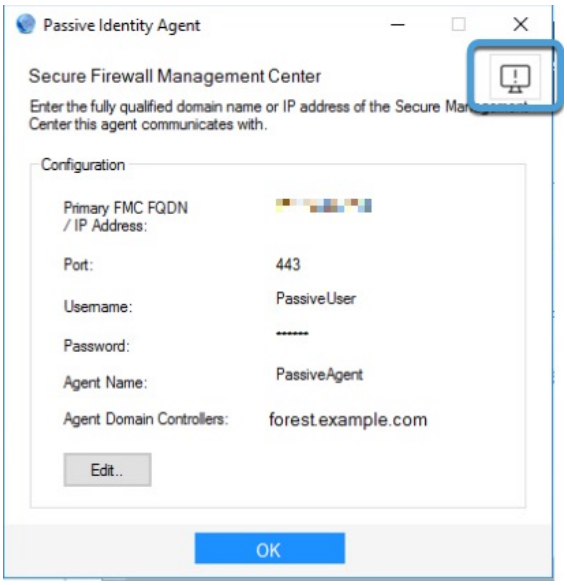
Do not rename **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent\CiscoPassiveIdentityAgentService.exe.config**; otherwise, the passive identity agent will stop generating log files. Do not remove or change the **.exe.config** file extension.

Step 2 Generate troubleshooting files by logging in to the Microsoft Active Directory domain controller.

Step 3 Start the passive identity agent software.

Step 4 Click the Troubleshooting button in the top right corner of the window.

The following figure shows an example.



The system displays a confirmation message and generates a .zip file containing troubleshooting files.

Your troubleshoot logs are saved to your system's Downloads folder; the file name starts with **TroubleshootLogs**.

Step 5 Manually view log files stored in plain text format in the agent's installation directory: **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent**.

Use Notepad or another text editor to view these files. Log files rotate after reaching 10MB in size.

Step 6 Use the Microsoft Active Directory event viewer to look for Kerberos-related events if you are not seeing user sessions in the Cloud-Delivered Firewall Management Center.

Look for the following Kerberos-related events:

- 4770
- 4768

For general information about audit policy, see Audit Policy Recommendations on learn.microsoft.com.

For more information about Windows Group Policy Object settings, see Group Policy Objects on learn.microsoft.com.

Security requirements for the passive identity agent

To safeguard the system, you should install the passive identity agent on a protected internal network. Although the passive identity agent is configured to have only the necessary services and ports available, you must make sure that attacks cannot reach it.

If the passive identity agent and the Security Cloud Control reside on the same network, you can connect the Security Cloud Control to the same protected internal network as the passive identity agent.

Regardless of how you deploy your appliances, inter-system communication is encrypted. You must ensure that communications between appliances are not interrupted, blocked, or tampered with; for example, with a distributed denial of service (DDoS) or man-in-the-middle attack.

Internet access requirements for the passive identity agent

By default, the passive identity agent is configured to communicate with the Firepower System over the internet using HTTPS on port 443/tcp (HTTPS). If you do not want the passive identity agent to have direct access to the internet, you can configure a proxy server.

If your Cloud-Delivered Firewall Management Center cannot communicate with the machine on which the passive identity agent is installed, you must use a proxy with the HTTPS protocol enabled.

You can choose how to set up the proxy. For example, you can use a commercial proxy and a Windows system proxy with HTTPS enabled to communicate with it.

This section lists the ports the passive identity agent uses to communicate with other agents, the Security Cloud Control, and Microsoft Active Directory.

Table 130: Passive Identity Agent port requirements

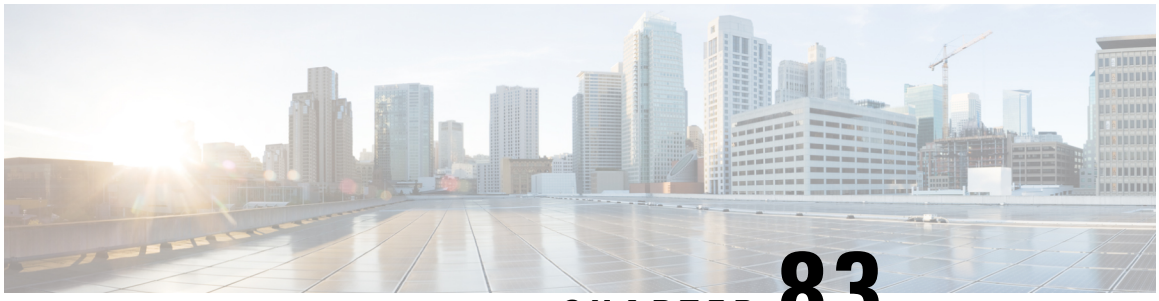
Port	Reason
443	Communicate with the Security Cloud Control.
135	Communicate with Microsoft Active Directory using the MSRPC protocol.
9095	Communicate with other agents using the UDP protocol.

History for the passive identity agent

This reference provides the version history and feature introduction timeline for the passive identity agent.

Table 131: History for the passive identity agent

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
Passive Identity Agent	November 8, 2024	7.1	This feature is introduced. Passive identity agent version 1.1 is compatible with 7.6.0 and later and adds the following: • You can use either FQDN, IPv4, or IPv6 to connect from the Passive Identity Agent to the Cloud-Delivered Firewall Management Center or Security Cloud Control. • Sends both IPv4 and IPv6 user sessions from Microsoft Active Directory (AD) to the Cloud-Delivered Firewall Management Center. • You can zip troubleshooting logs. • When you start the passive identity agent software, a list of prerequisites is displayed. The passive identity agent identity source sends session data from Microsoft Active Directory (AD) to the Cloud-Delivered Firewall Management Center. Passive identity agent software is supported on: • Microsoft AD server (Windows Server 2008 or later) • Microsoft AD domain controller (Windows Server 2008 or later) • Any client connected to the domain you want to monitor (Windows 8 or later) See: User Control With the Passive Identity Agent.



CHAPTER 83

Identity Source: pxGrid Cloud Identity (ISE 3.3 and Earlier)

This document discusses how to configure and use the pxGrid Cloud Identity Source with Cisco ISE version 3.3 and earlier.

- pxGrid Cloud identity source, on page 2433
- pxGrid Cloud identity source configuration, on page 2436
- Enable the pxGrid Cloud service in Cisco ISE, on page 2439
- Register Cisco ISE with the Catalyst Cloud Portal, on page 2439
- Register the pxGrid Cloud connection with Cisco ISE, on page 2442
- Create a pxGrid Cloud identity source, on page 2443
- Create dynamic attributes filters, on page 2455
- Create access control rules or DNS rules using dynamic attributes filters, on page 2456
- Troubleshoot the pxGrid Cloud identity source, on page 2457
- Deactivating and deleting the pxGrid Cloud identity source, on page 2459

pxGrid Cloud identity source

The Cisco Identity Services Engine (Cisco ISE) pxGrid Cloud identity source enables you to use subscription and user data from Cisco ISE in Cloud-Delivered Firewall Management Center access control rules. Also, the identity source uses constantly changing dynamic objects from Cisco ISE in access control policies in the Cloud-Delivered Firewall Management Center.

Integration technologies

The pxGrid Cloud identity source also uses:

- The Cisco Platform Exchange Grid (pxGrid), which enables multivendor, cross-platform network system collaboration in things like security monitoring and detection systems, network policy platforms, asset and configuration management, identity, and access management. pxGrid Cloud is the cloud-based interface to Cisco ISE.

  More information about pxGrid is available in resources such as What is PxGrid? on devnet.

- The Cisco Digital Network Architecture (Cisco DNA) delivers automation, security, predictive monitoring, and a policy-driven approach. It provides end-to-end network visibility and uses network insights to optimize network performance and deliver the best user and application experience.

To use the pxGrid Cloud identity source with the Security Cloud Control, you must Create a Cisco Account.

- What is pxGrid? on devnet
- Cisco Platform Exchange Grid Cloud on devnet

Related Topics

Configuring a pxGrid Cloud identity source, on page 2468
Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436

Restrictions for the pxGrid Cloud identity source

Before you set up the pxGrid Cloud identity source, note these restrictions:

- pxGrid Cloud suppports these regions: `us-west-2`, `eu-central-1`, and `ap-southeast-1`.

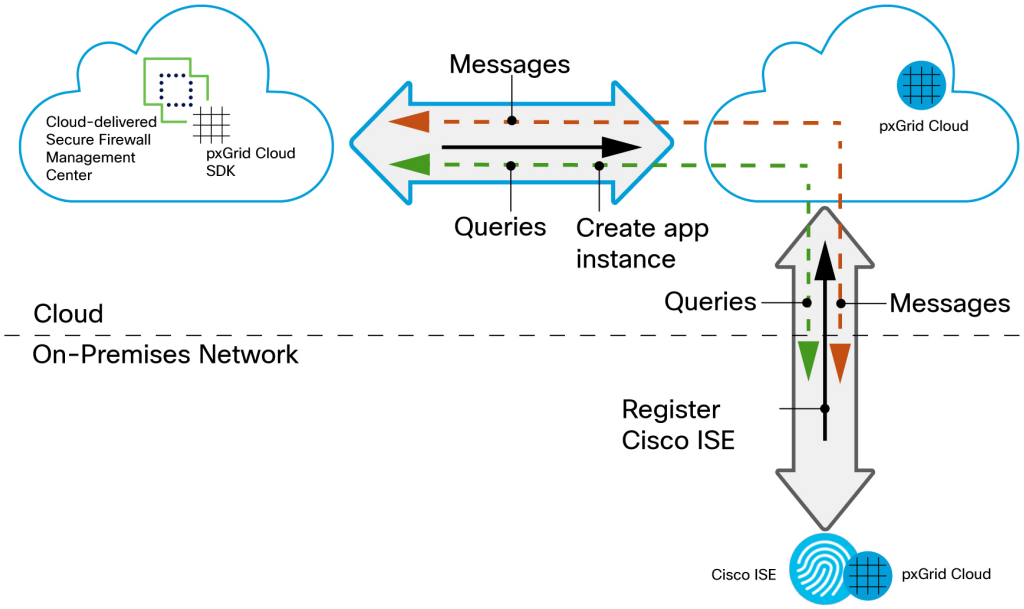
How the pxGrid Cloud identity source works

Summary

The key components involved in the pxGrid Cloud identity source are:

- Cloud-Delivered Firewall Management Center: Uses the pxGrid Cloud SDK to programmatically retrieve user information
- Cisco ISE: Provides user information, SGT, endpoint profile, and other details
- pxGrid Cloud: Facilitates secure data exchange between the management center and ISE
- Authentication process: Requires one-time passwords (OTP) to establish trusted communication

Workflow



These stages describe how the pxGrid Cloud identity source works:

1. In Cisco ISE, the administrator enables the use of pxGrid Cloud.
2. The administrator registers Cisco ISE as a product in pxGrid Cloud, which authenticates Cisco ISE and pxGrid Cloud and enables them to communicate with each other.

   The authentication process requires you to paste a one-time password (OTP) from pxGrid Cloud into Cisco ISE.
3. In pxGrid Cloud, the administrator creates an "app instance" that generates an OTP for use in the Cloud-Delivered Firewall Management Center to authenticate the two with each other.
4. After completing all the preceding tasks, the Cloud-Delivered Firewall Management Center (which includes the pxGrid Cloud SDK) can query Cisco ISE using pxGrid Cloud and retrieve sessions containing user information, SGT, endpoint profile, and other details.
5. Many types of dynamic objects can be filtered and sent to the Cloud-Delivered Firewall Management Center as dynamic objects to be used in access control rules. These include: SGT, endpoint profile, posture status, and machine authentication.

   User information is retrieved from Cisco ISE and group information is retrieved from either Microsoft Active Directory or Azure Active Directory.

Related Topics

Configuring a pxGrid Cloud identity source, on page 2468
Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436

pxGrid Cloud identity source configuration

These topics summarize how to configure a pxGrid Cloud identity source either for ISE 3.3 and earlier or for ISE 3.4 and later. The steps are different so make sure you follow them exactly.

Related Topics

Configuring a pxGrid Cloud identity source, on page 2468
Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436
Configuring a pxGrid Cloud identity source, on page 2468
Enable the pxGrid Cloud service in Cisco ISE, on page 2470
Create an app instance, on page 2473
Create the identity source, on page 2445
Activate the app instance, on page 2446
Activate the pxGrid Cloud identity source, on page 2448

Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier)

Before you begin, create a Cisco Account.



Important *This topic applies to Cisco ISE version 3.3 or earlier*. If you are using a later version, see Configuring a pxGrid Cloud identity source, on page 2468 instead.

Summary

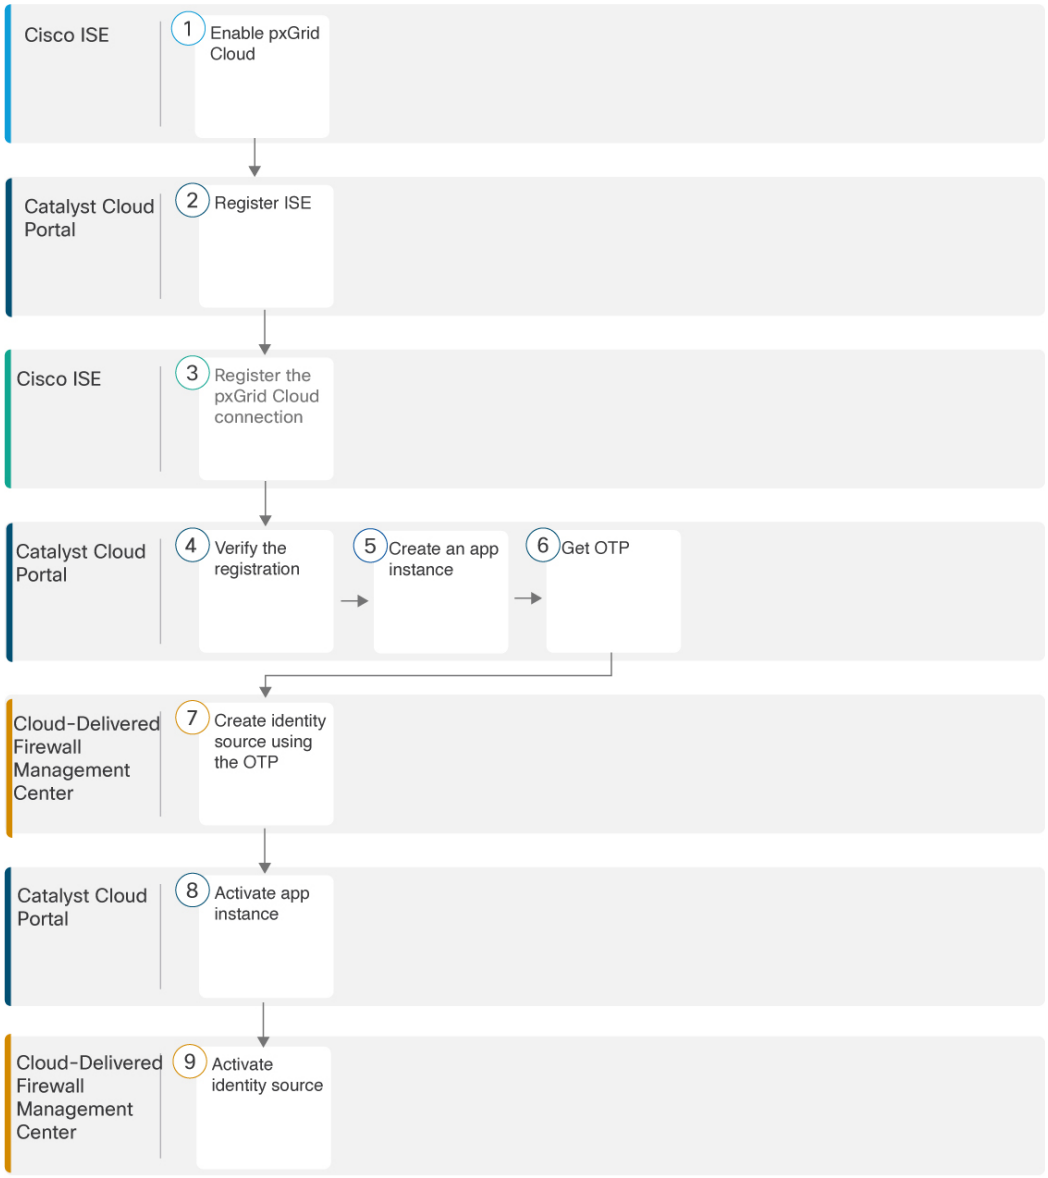
The key components involved in configuring a pxGrid Cloud identity source are:

- Cisco ISE: Enables pxGrid Cloud for secure data exchange and provides the data source for user information
- Catalyst Cloud Portal: Manages registration and authentication between components, creates application instances, and provides one-time passwords for secure connection
- Cloud-Delivered Firewall Management Center: Creates and activates the identity source connection to receive user data from Cisco ISE

This process establishes a secure cloud-based connection that enables the Cloud-Delivered Firewall Management Center to receive user and group information from Cisco ISE through the pxGrid Cloud service.

Workflow

Figure 390: Configure a pxGrid Cloud identity source



Lorem ipsum

1 Enable the pxGrid Cloud service in Cisco ISE, on page 2439
2 Register Cisco ISE with the Catalyst Cloud Portal, on page 2439
3 Register Cisco ISE with the Catalyst Cloud Portal, on page 2439
4 Create an app instance, on page 2443
5 Create the identity source, on page 2445

6 Activate the app instance, on page 2446
7 Activate the app instance, on page 2446

These stages describe configuring a pxGrid Cloud identity source:

1. Cisco ISE enables pxGrid Cloud functionality.
   - pxGrid Cloud enables you to subscribe to offers and to register apps (in this case, the Cloud-Delivered Firewall Management Center) for secure data exchange in a cloud environment

   For more information, see Enable the pxGrid Cloud service in Cisco ISE, on page 2439.
2. The administrator registers Cisco ISE in the Catalyst Cloud Portal and authenticates communication between Cisco ISE and the Catalyst Cloud Portal.

   For more information, see Register Cisco ISE with the Catalyst Cloud Portal, on page 2439.
3. The administrator registers the pxGrid Cloud with Cisco ISE and verifies the registration.

   For more information, see Register the pxGrid Cloud connection with Cisco ISE, on page 2442.
4. The administrator creates an application instance in the Catalyst Cloud Portal and gets the one-time password (OTP).
   - The application instance enables the Cloud-Delivered Firewall Management Center to authenticate with Cisco ISE using the pxGrid Cloud service
   - The OTP, required for the next step, expires in 60 minutes
5. The administrator creates the pxGrid Cloud identity source using the OTP from the previous step.
   - Linking the app enables the Cloud-Delivered Firewall Management Center to authenticate with Cisco ISE and the Catalyst Cloud Portal so it can receive user data from Cisco ISE

   For more information, see Create the identity source, on page 2445.
6. The administrator activates the app instance in the Catalyst Cloud Portal.

   For more information, see Activate the app instance, on page 2446.
7. The administrator activates the pxGrid Cloud identity source in the Cloud-Delivered Firewall Management Center.

   For more information, see Activate the pxGrid Cloud identity source, on page 2448

What's next

After you have completed all the preceding tasks, you can:

- Test the pxGrid Cloud identity source to make sure it's working properly. For more information, see Test the pxGrid Cloud identity source, on page 2450.
- Create dynamic attributes filters, which define what dynamic objects are sent to the Cloud-Delivered Firewall Management Center. For more information, see Create dynamic attributes filters, on page 2455.
- After you configure the pxGrid Cloud identity source, you can use dynamic objects, Microsoft AD user and groups, and Azure AD users and groups in access control rules.

Enable the pxGrid Cloud service in Cisco ISE

Enable the pxGrid Cloud service to establish outbound connections to Cisco pxGrid Cloud for cloud-based integration capabilities.

The pxGrid Cloud service allows Cisco ISE to connect to cloud-based services and share identity information securely. This service can be enabled on multiple nodes for high availability.

Before you begin

- Ensure that you install and activate the Advantage license tier in your Cisco ISE deployment.
- The pxGrid Cloud agent creates an outbound HTTPS connection to Cisco pxGrid Cloud. Therefore, you must configure Cisco ISE proxy settings if the customer network uses a proxy to reach the internet. To configure proxy settings in Cisco ISE, click the **Menu** icon (☰) and choose **Administration > System > Settings > Proxy**.
- The Cisco ISE Trusted Certificates Store must include the root CA certificate required to validate the server certificate presented by pxGrid Cloud. Ensure that the **Trust for Authentication of Cisco Services** option is enabled for this root CA certificate. To enable **Trust for Authentication of Cisco Services**, navigate to **Administration** > **System** > **Certificates**.

Follow these steps to enable the pxGrid Cloud service:

Procedure

Step 1 In the Cisco ISE GUI, click the **Menu** icon (☰) and choose **Administration > System > Deployment**.

Step 2 Click the node on which you want to enable the pxGrid Cloud service.

Step 3 In the **General Settings** tab, enable the **pxGrid** service.

Step 4 Check the **Enable pxGrid Cloud** check box.

The pxGrid Cloud service can be enabled on two nodes to enable high availability.

Note
You can enable the **pxGrid Cloud** option only when the **pxGrid** service is enabled on that node.

Register Cisco ISE with the Catalyst Cloud Portal

Register Cisco ISE as an app in the Catalyst Cloud Portal to authenticate communication between the systems.

This task discusses how to register Cisco ISE as an app in the Catalyst Cloud Portal and to authenticate communication between the Catalyst Cloud Portal and Cisco ISE.

Also refer to Register Cisco ISE in the *pxGrid Cloud Solution Guide*.

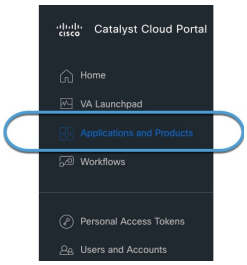
Procedure

Step 1 Log in to the Cisco Cloud Catalyst Portal.

If prompted, choose an account to use.

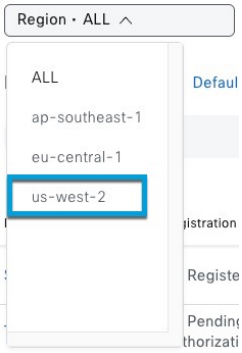
Step 2 Click **Register**.

Step 3 In the Catalyst Cloud Portal, click ☰ > **Applications and Products**.



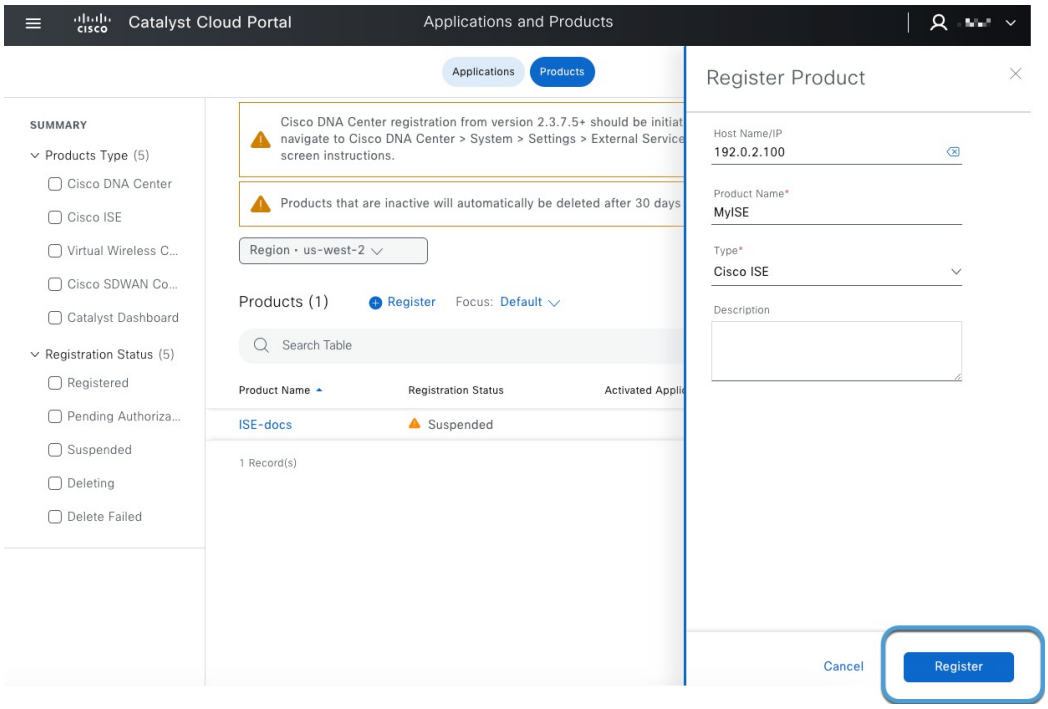
Step 4 At the top of the page, click **Products**.

Step 5 From the **Region** list, click **us-west-2**, **eu-central-1**, or **ap-southeast-1**.



Step 6 Click **Register**.

The registration page appears.



Step 7 Enter the required information.

- **Host Name/IP**: (Optional.) Enter the ISE server's fully qualified domain name or IP address. If you enter an IP address, omit the scheme (for example, **https://**) and the port, if any.
- **Product Name**: Enter a unique name to identify this server.
- **Type**: From the list, click **Cisco ISE**.
- **Description**: Enter an optional description.

Step 8 Click **Register**.

Step 9 Generate a one-time password (OTP).

- If you've previously registered ISE apps and see yours listed, click **Generate OTP** in the **Actions** column; you'll need it in the next part of this procedure.
- If you're registering your app now, the OTP is displayed. Click to copy it to the clipboard; you'll need it in the next part of this procedure.

What to do next

See Register the pxGrid Cloud connection with Cisco ISE, on page 2442.

Register the pxGrid Cloud connection with Cisco ISE

This task registers the pxGrid Cloud connection with Cisco ISE to enable pxGrid Cloud to send user data to the pxGrid Cloud identity source in Security Cloud Control.

Registering this connection enables communication between your identity services and cloud security services, facilitating the exchange of user authentication and authorization data.

Before you begin

Complete the tasks discussed in Register Cisco ISE with the Catalyst Cloud Portal, on page 2439.

Follow these steps to register the pxGrid Cloud connection with Cisco ISE:

Procedure

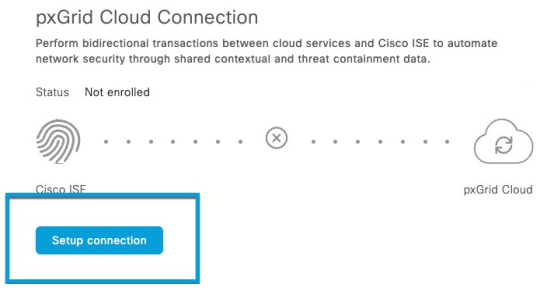
Step 1 Log in to Cisco ISE as an administrator.

Step 2 Click ☰ > **Administration** > **pxGrid Services** > **Client Management** > **pxGrid Cloud Connection**.

Step 3 Make sure all services are enabled with read/write privileges.

Step 4 In the left navigation bar, click **pxGrid Cloud Connection**.

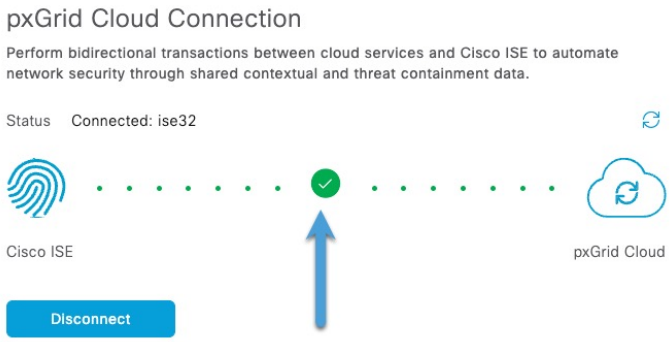
Step 5 Click **Setup Connection**.



Step 6 Paste the OTP value in the provided field.

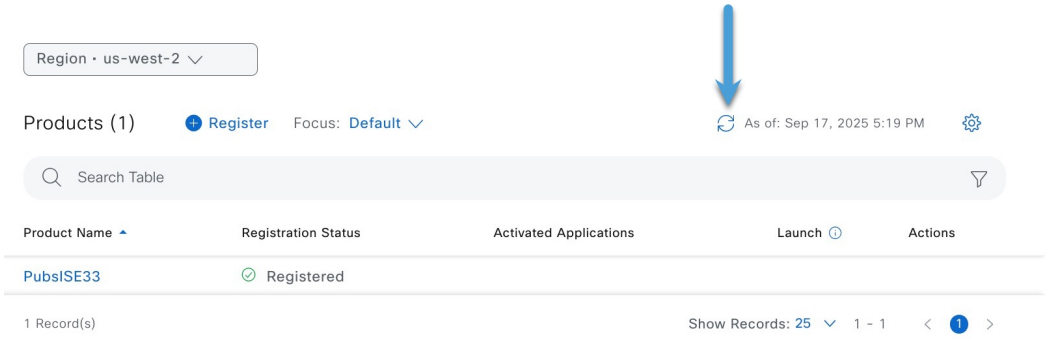
Step 7 Click **Connect**.

A green check mark confirms that connection was successful.



Step 8 Confirm the setup has been successful:

a) Log in to the Catalyst Cloud Portal.
b) Click the **Products** tab.
c) Click **Refresh**.



d) Verify that **Registered** is displayed as the status of your product.

What to do next

Continue with Create the identity source, on page 2445.

Create a pxGrid Cloud identity source

The tasks discuss how to create a pxGrid Cloud identity source using Cisco ISE, the Catalyst Cloud Portal, and Security Cloud Control. You must complete all tasks in the order shown; in some cases, there is a time limit due to the expiration of a required One-Time Password (OTP).

Related Topics

Create an app instance, on page 2443
Create the identity source, on page 2445
Activate the app instance, on page 2446
Activate the pxGrid Cloud identity source, on page 2448
Test the pxGrid Cloud identity source, on page 2450

Create an app instance

Create an application instance in the Catalyst Cloud Portal to establish a pxGrid Cloud identity source.

This task is one of several tasks you must perform to create a pxGrid Cloud identity source to send user session data to the Cloud-Delivered Firewall Management Center.

There is a one-hour time limit on the one-time password (OTP) required to complete this procedure successfully. You do not need to log in to Cisco ISE.

Before you begin

Complete all of these tasks first:

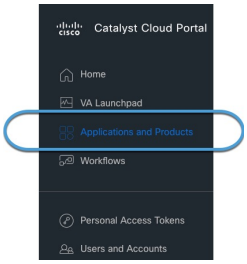
- Enable the pxGrid Cloud service in Cisco ISE, on page 2439

- Register Cisco ISE with the Catalyst Cloud Portal, on page 2439
- Register the pxGrid Cloud connection with Cisco ISE, on page 2442

Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

Step 2 In the Catalyst Cloud Portal, click ☰ > **Applications and Products**.



Step 3 At the top of the page, click **Applications**.

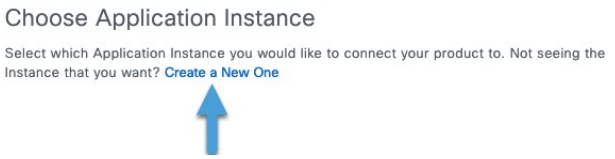
Step 4 From the **Regions** list, click **us-west-2**, **eu-central-1**, or **ap-southeast-1**.

Step 5 Click **Manage** (or **Activate**) next to **Firepower Management Center**.

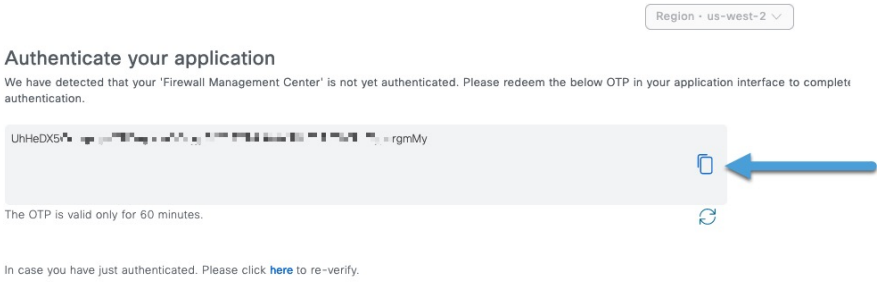
Step 6 Click **Add**.

Step 7 Click **Create a New One**.

The following figure shows an example.



Step 8 Click the copy button next to the displayed OTP.



Copy the OTP to a text file; it expires in 60 minutes.

Step 9 Continue with Create the identity source, on page 2445.

Create the identity source

This task creates a pxGrid Cloud identity source to send user session data to the Cloud-Delivered Firewall Management Center.

This task is one of several required to create a pxGrid Cloud identity source to send user session data to the Cloud-Delivered Firewall Management Center.

Before you begin

Complete the task discussed in Create an app instance, on page 2443.

Procedure

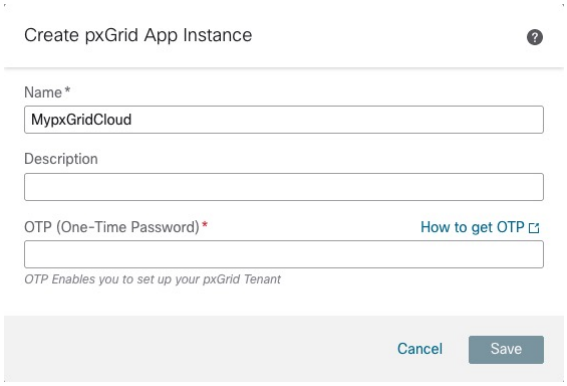
Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Step 2 Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**

Step 3 Click **Identity Services Engine (pxGrid Cloud)**.

Step 4 Click **Create pxGrid Application Instance**.

This figure shows an example.



Step 5 Enter the required information.

Value	Description
Name	Enter a name to uniquely identify this connector.
Description	Optional description.
OTP (One-Time Password)	Enter the OTP.

Step 6 Click **Create**.

Step 7 At the top of the page, click **Save**.

What to do next

Continue with Activate the app instance, on page 2446.

Activate the app instance

Activate an app instance to create a pxGrid Cloud identity source that sends user session data to the Cloud-Delivered Firewall Management Center.

There is a one-hour time limit on the one-time password (OTP) required to complete this procedure successfully. You do not need to log in to Cisco ISE.

Before you begin

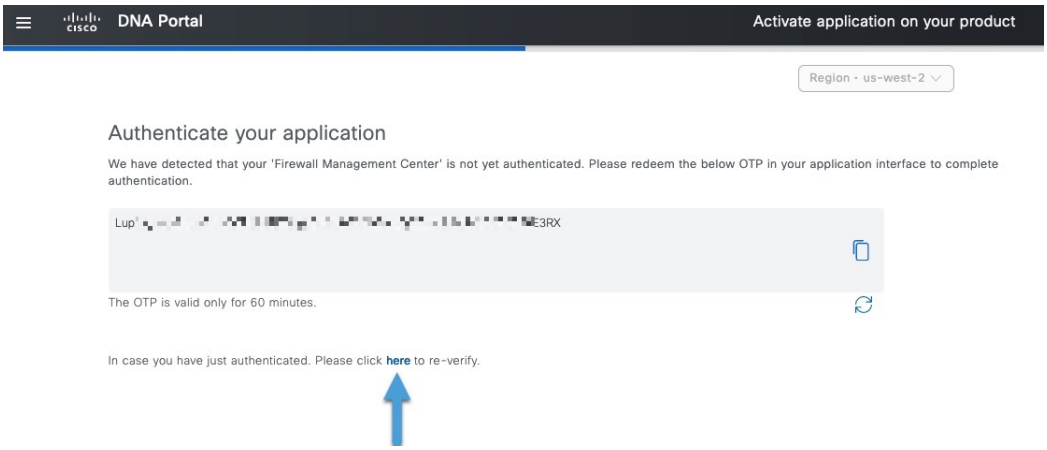
Complete the task discussed in Create the identity source, on page 2445.

Follow these steps to activate the app instance:

Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

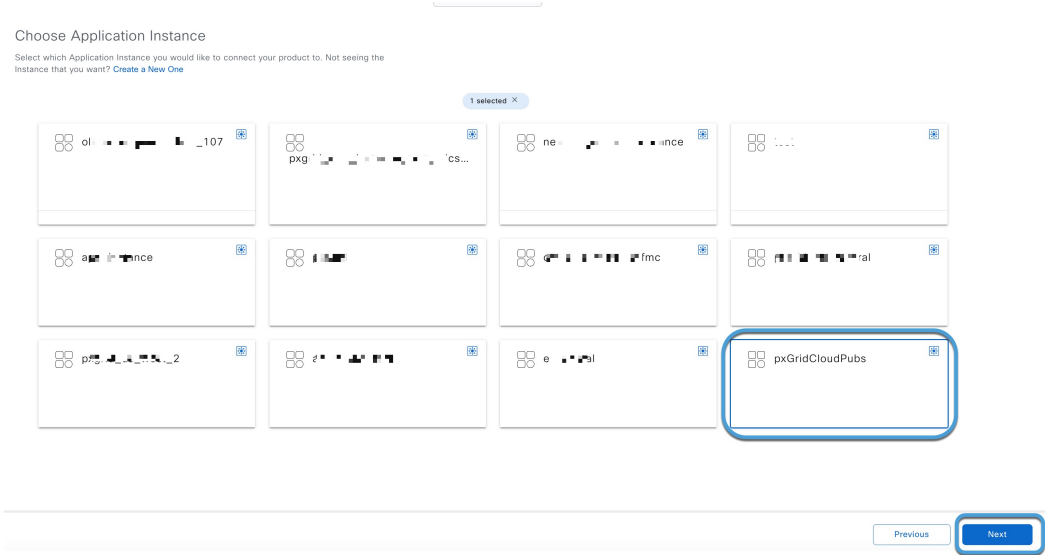
Step 2 Reverify the app by clicking the word **here** as the following figure shows.



Step 3 Click the name of the application instance you just created in Cloud-Delivered Firewall Management Center.

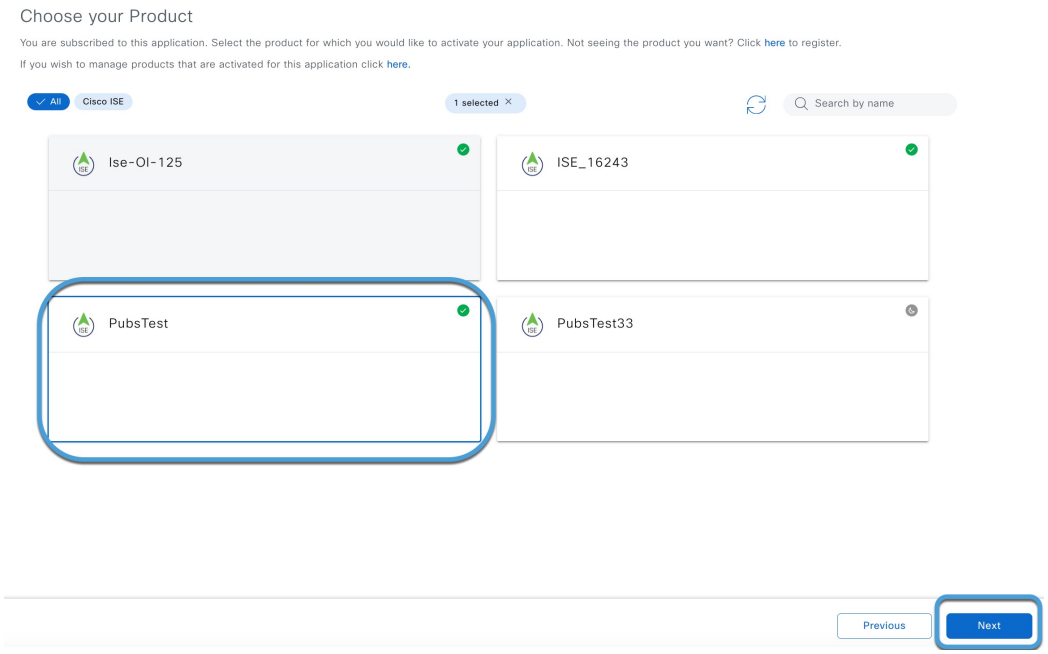
Step 4 Click **Next**.

Example:

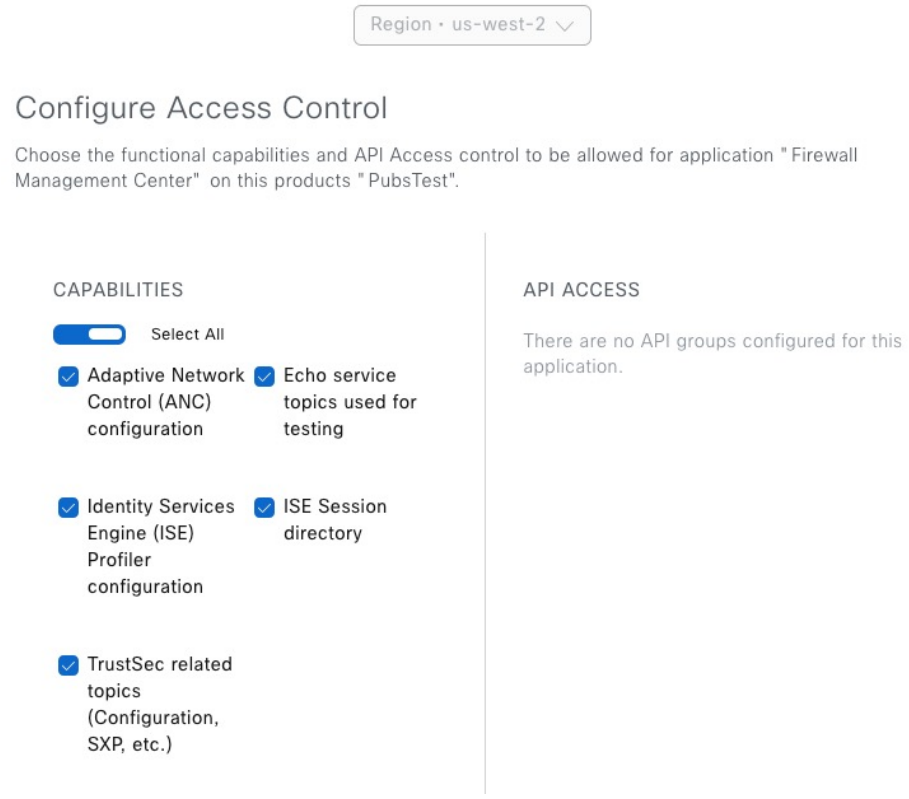


Step 5 On the Choose Product page, click the name of the Cisco ISE product and click **Next**.

Example:



Step 6 Select the check box next to each scope.
The following figure shows an example.



Step 7 Click **Next**.

Step 8 Review the displayed information for accuracy. Make sure all scopes are selected.

Step 9 Click **Activate**.

It can take several minutes for the app instance to be activated.

Continue with Activate the pxGrid Cloud identity source, on page 2448.

Activate the pxGrid Cloud identity source

This task activates the pxGrid Cloud identity source to enable identity services integration in your security infrastructure.

This task explains how to activate the pxGrid Cloud identity source in the Security Cloud Control.

Before you begin

Complete the tasks discussed in Activate the app instance, on page 2446.

Note Only one pxGrid Cloud identity source can be active at a time.

Procedure

Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

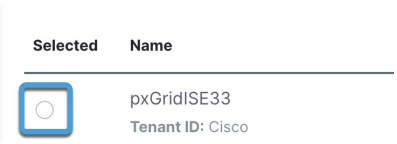
Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 2 Click **Identity Services Engine (pxGrid Cloud)**.

Step 3 Click **Save** at the top of the page.

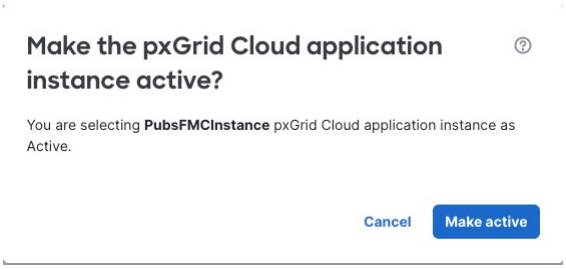
Step 4 If a green check mark is *not* displayed next to the name of the identity source, select it.

Example:



Step 5 Click **Make Active**.

Example:



Step 6 (Optional.) Select the following options if desired:

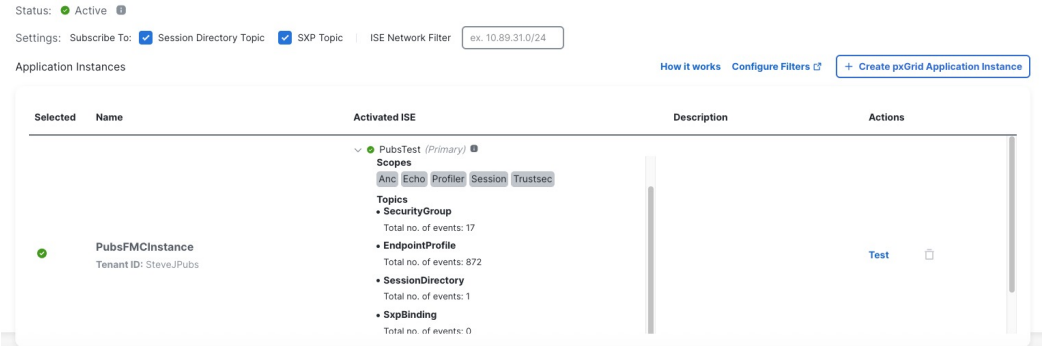
- **Session Directory Topic**: Select the check box to receive ISE user session information from the Cisco ISE server.
- **SXP Topic**: Select the check box to receive updates to SGT-to-IP mappings when available from the ISE server. This option is required to use destination SGT tagging in access control rules.
- **ISE Network Filter**: Optional filter you can set to restrict the data that Cisco ISE reports. If you provide a network filter, Cisco ISE reports data from the networks in that filter.

  You have the following options:

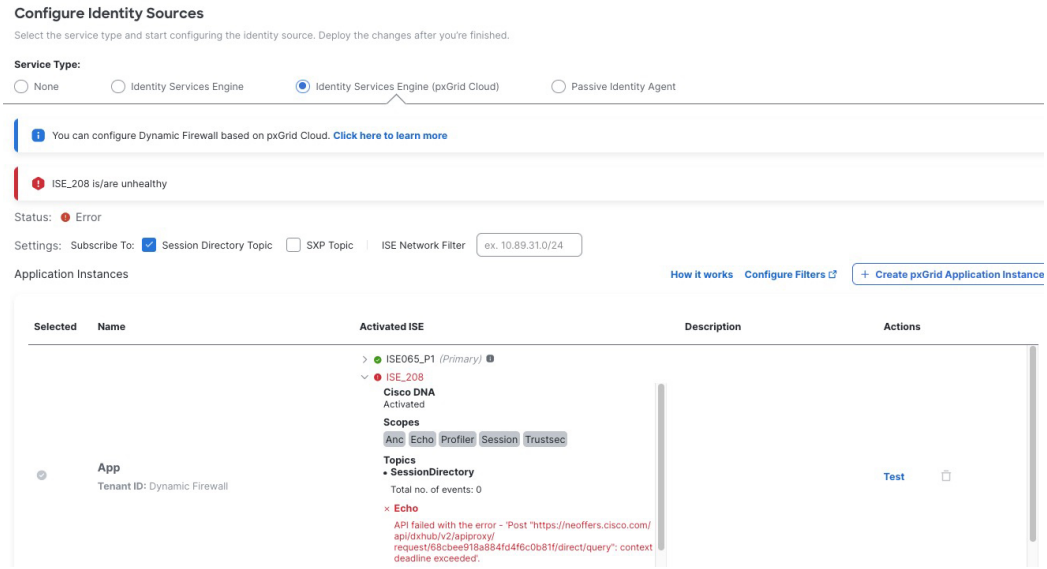
  - Leave the field blank to specify **`any`**.
  - Enter a single IPv4 address block using CIDR notation.
  - Enter a list of IPv4 address blocks using CIDR notation, separated by commas.

Step 7 Under Activated ISE, expand the identity source.

Example *normal* result:



Example *error* result:



In the event of an error, see Test the pxGrid Cloud identity source, on page 2450.

Step 8 Verify the status is Active and that all scopes and topics are displayed.

Wait a few minutes for data to be downloaded.

What to do next

See Test the pxGrid Cloud identity source, on page 2450.

Test the pxGrid Cloud identity source

This task enables you to perform diagnostics using the Security Cloud Control to determine if the identity source is working.

Errors might include communication with Cisco ISE, or with the Cisco ISE configuration with Catalyst Cloud Portal.

Procedure

Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Step 2

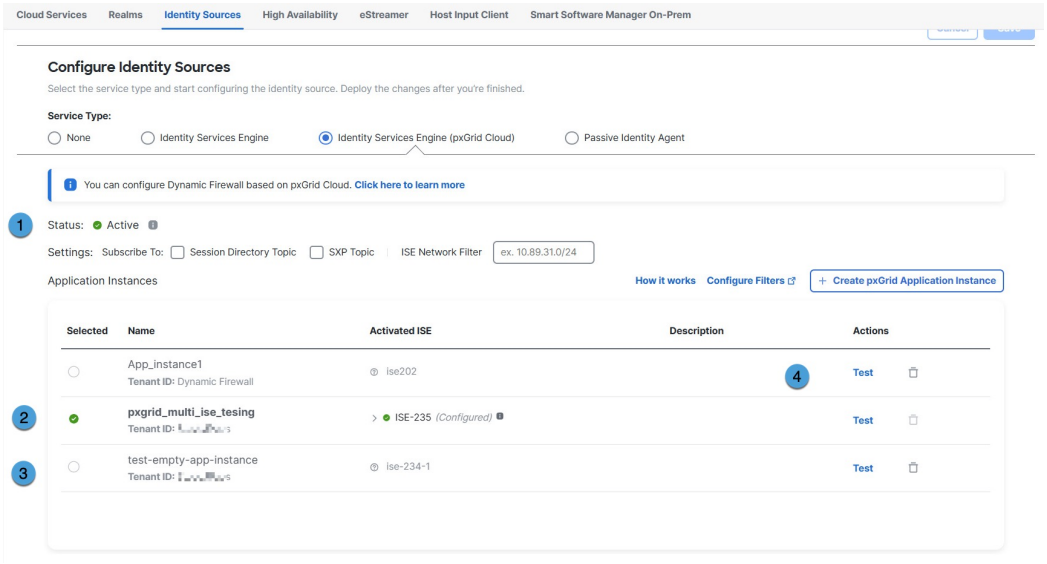
Step 3 Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 4

Step 5 Click **Identity Services Engine (pxGrid Cloud)**.

Step 6 Review the configuration status information.

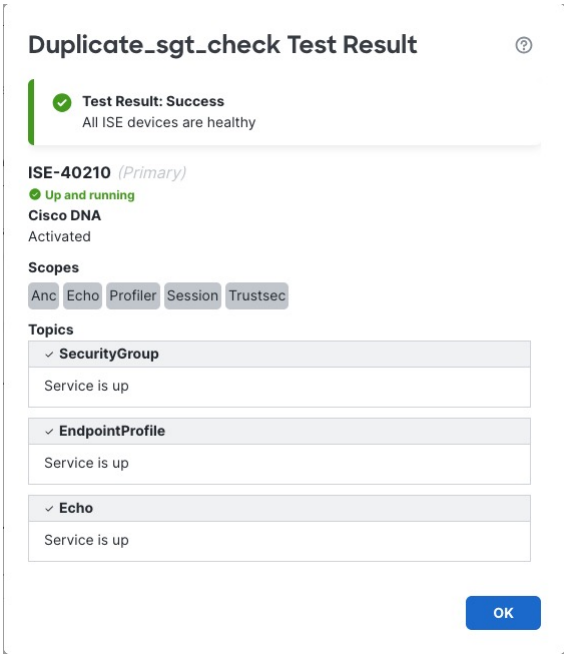
The following figure shows an example configuration.



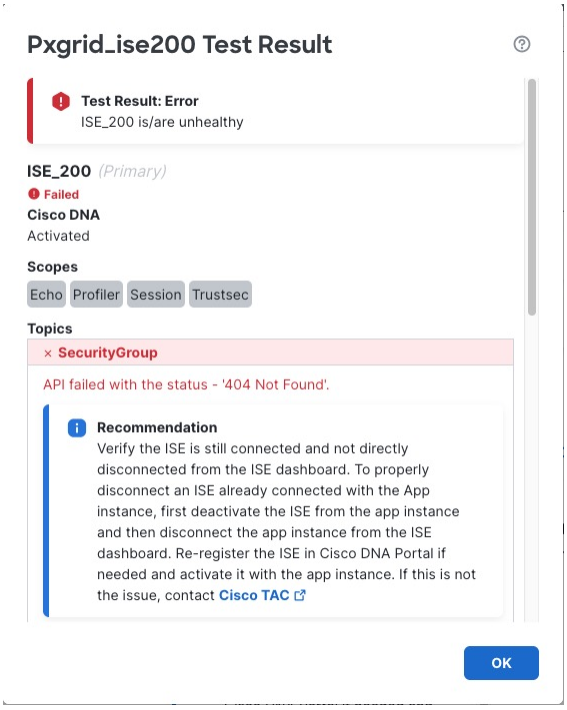
The following table has more information about the numbered areas in the figure.

Number	Meaning
1	Overall status Any errors in the overall status of the Cisco ISE app instances are displayed. In that case, scroll to that instance and either expand the error message or click **Test** for more information.
2	Active A green check mark indicates the app is active.
3	Inactive A dimmed app instance is inactive. You can activate it by selecting the check box next to its name and then clicking **Make active**.
4	**Test** button Click **Test** to perform diagnostic tests that show more detailed status of the app instance.

The following figure shows a sample success message.



The following figure shows an example error result.



Step 7 Click **Test** to perform diagnostic tests.

If errors occur, use the following error code reference to help diagnose and solve issues with Cisco ISE, pxGrid Cloud, and the Catalyst Cloud Portal. If these suggestions do not work, or if you have a different issue, contact Cisco TAC.

Table 132: Error code troubleshooting

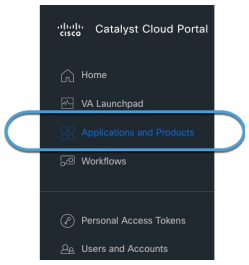
Error Code	Troubleshooting Steps
403 – Forbidden	Verify the Cisco ISE product is not in a **Pending** or **Suspended** state in the Catalyst Cloud Portal. If suspended, verify that Cisco ISE is registered as discussed in Enable pxGrid Cloud service in Cisco ISE and register your device. Additionally, verify pxGrid Cloud services are publicly available.
404 – Not Found	Verify the Cisco ISE server is not directly disconnected from the Cisco ISE dashboard. To properly disconnect Cisco ISE already connected with the app instance, first deactivate Cisco ISE from the app instance and then disconnect the app instance from the Cisco ISE dashboard.
408 – Request Timeout	Check whether there are any general connectivity issues with Cisco ISE and verify pxGrid Cloud connectivity status is **Connected** in the ISE dashboard under ☰ > **Administration** > **pxGrid Services** > **Client Management** > **pxGrid Cloud Connection**. Verify the Cisco ISE server is not directly disconnected from the Cisco ISE dashboard.
413 – Content Too Large	Review the pxGrid Cloud API limitations on GitHub. If needed, consider upgrading your Cisco ISE version to fully utilize pxGrid Cloud support.
500 – Internal Server Error	Check that the Cisco ISE server is operational and that pxGrid Cloud services are active (verify MNT, SXP, pxGrid nodes, and so on). For more information, see Monitoring and debugging in the Cisco pxGrid chapter in the *Cisco Identity Services Engine Administrator Guide*.

Step 8 If the product is in a **Pending** or **Suspended** state, verify that it is active.

a) Log in to the Catalyst Cloud Portal.

b) In the Catalyst Cloud Portal, go to ☰ > **Applications and Products**.

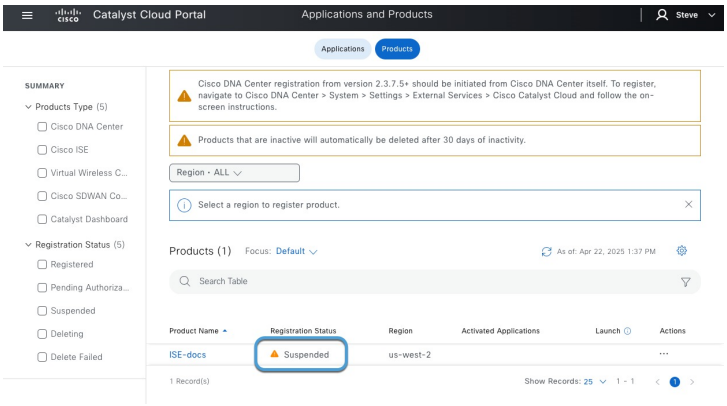
Example:



c) Click the **Products** tab.

Example:

The following figure shows an example of a suspended product.

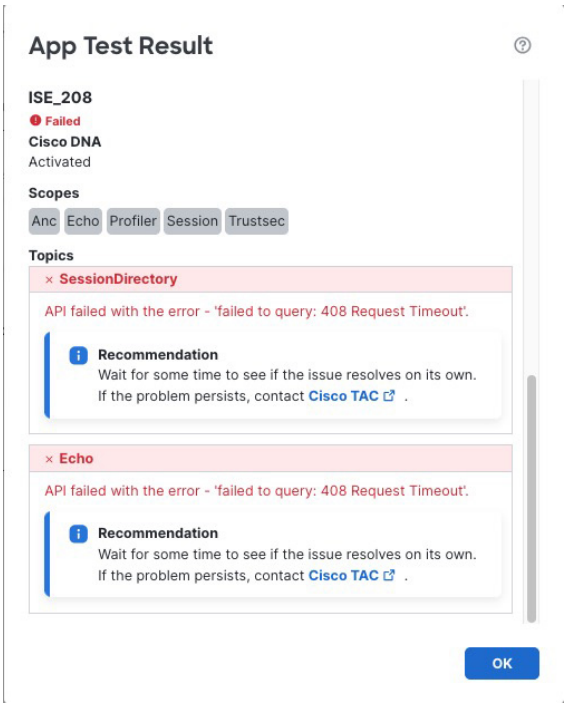


d) To correct the issue, in the Actions column, click ••• and click **Generate OTP**.

e) Use the OTP as discussed in Create the identity source, on page 2445.

Step 9 If you encounter a cluster member not reachable error, find what node is not reachable.

If a member of the Cisco ISE cluster is not reachable, a page like this is displayed:



To find what node is not reachable, log in to Cisco ISE primary administration node as an administrator and click click the **Menu** icon (☰) and choose **Administration** > **System** > **Deployment**, then see Node Status in a Cisco ISE Deployment.

Create dynamic attributes filters

Create dynamic attributes filters to expose them as dynamic objects in Security Cloud Control for use in access control policies.

Dynamic attributes filters that you define using the Dynamic Attributes Connector are exposed in the Security Cloud Control as dynamic objects that can be used in access control policies. For example, restrict access to an AWS server for the Finance Department to only members of the Finance group defined in Microsoft Active Directory.

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

For more information about access control rules, see Create access control rules or DNS rules using dynamic attributes filters, on page 2456.

Before you begin

Connectors, on page 2631

Procedure

Step 1 Log in to Security Cloud Control.

a) Click **Firewall**.

b) Click **Administration** > **Dynamic Attributes Connector** > **Dynamic Attributes Filters**.

Step 2 Do any of the following:

- Add a new filter: click **Add** ( ).
- Edit or delete a filter: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 3 Enter the following information.

Item	Description
Name	Unique name to identify the dynamic filter (as a dynamic object) in a policy and in the Security Cloud Control Object Manager (**External Attributes** > **Dynamic Object**).
Connector	From the list, click the name of a connector to use.
Query	Click Add .

Step 4 To add or edit a query, enter the following information.

Item	Description
Key	Click a key from the list. Keys are fetched from the connector.
Operation	Click one of the following: • **Equals** to exactly match the key to the value. • **Contains** to match the key to the value if any part of the value matches.
Values	Click either **Any** or **All** and click one or more values from the list. Click **Add another value** to add values to your query.

Step 5 Click **Show Preview** to display a list of networks or IP addresses returned by your query.

Step 6 When you're finished, click **Save**.

Step 7 (Optional.) Verify the dynamic object in the Security Cloud Control.

a) Log in to the Security Cloud Control.

b) Click **Policies** > **Firewall Threat Defense**.

c) Click **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.

The dynamic attribute query you created should be displayed as a dynamic object.

Create access control rules or DNS rules using dynamic attributes filters

Create access control rules using dynamic objects that are based on dynamic attributes filters you have previously configured.

This topic discusses how to create access control rules using dynamic objects (these dynamic objects are named after the dynamic attributes filters you created previously).

To add dynamic attributes filters to DNS policies, see Create basic DNS policies, on page 2561.

To add dynamic attributes filters to DNS policies, see Creating Basic DNS Policies.

Before you begin

Create dynamic attributes filters as discussed in Create dynamic attributes filters, on page 2455.

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

Follow these steps to create access control rules or DNS rules using dynamic attributes filters:

Procedure

Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Click **Policies** > **Firewall Threat Defense** > **Access Control heading** > **Access Control**.

Step 2 Click **Firewall**.

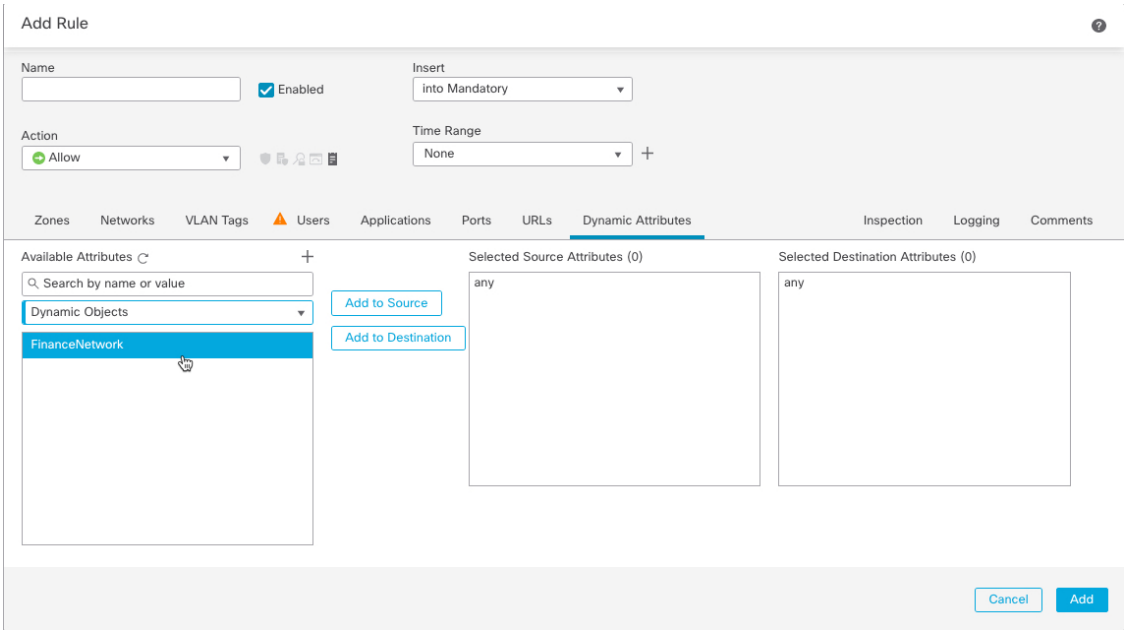
Step 3 Click **Edit** ( ) next to an access control policy.

Step 4 Click **Add Rule**.

Step 5 Click the **Dynamic Attributes** tab.

Step 6 In the Available Attributes section, from the list, click **Dynamic Objects**.

The following figure shows an example.



This example shows a dynamic object named `APIC Dynamic Attribute` that corresponds to the dynamic attribute filter created in the dynamic attributes connector.

Step 7 Add the desired object to source or destination attributes.

Step 8 Add other conditions to the rule if desired.

See Dynamic attributes rule conditions, on page 2681.

Troubleshoot the pxGrid Cloud identity source

These topics describe troubleshooting the pxGrid Cloud identity source.

Related Topics

Primary device cannot be processed, on page 2458
Primary device cannot be processed, on page 2458

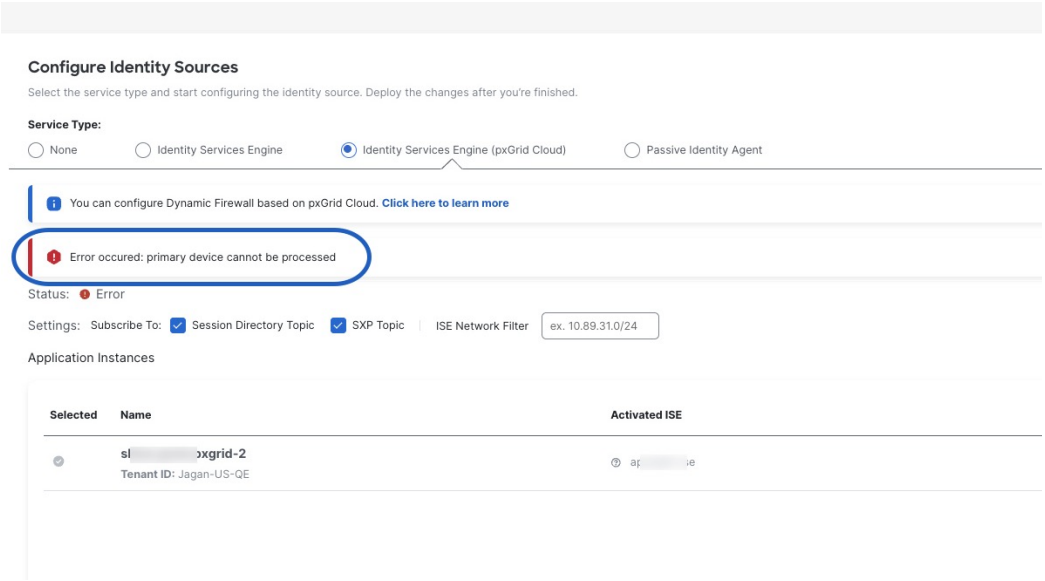
Primary device cannot be processed

This task resolves the error "primary device cannot be processed" or "ISE is unhealthy" that occurs when a Cisco ISE cluster is associated with more than one app instance.

Each Cisco ISE cluster must be associated with one and only one app instance, typically in a single dedicated tenant.

If you associate a Cisco ISE with more than one app instance, an error such as `Error occured: primary device cannot be processed` or `ISE is unhealthy` is displayed for the identity source.

Example:



The solution is to deactivate the ISE properly from other app instances for that tenant, before using it in any other tenant or app instance

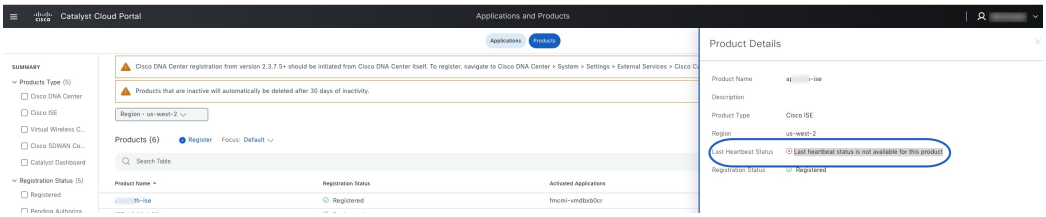
Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

Step 2 At the top of the page, click **Applications**.

Step 3 Locate a Cisco ISE product that is activated and verify it is the one causing the issue.

Example:



Step 4 Wait for the product to be removed.

Step 5 Deactivate the app instance as described in Deactivate the pxGrid Cloud app instance, on page 2459.

Deactivating and deleting the pxGrid Cloud identity source

Deactivating the FMC app instance in the Catalyst Cloud Portal is an optional task to troubleshoot issues with the Cisco ISE integration.

You should delete the pxGrid Cloud identity source from the Security Cloud Control when you are certain you don't want to use it again.

Related Topics

Deactivate the pxGrid Cloud app instance, on page 2459
Delete the pxGrid Cloud identity source, on page 2462

Deactivate the pxGrid Cloud app instance

This task explains how to deactivate a pxGrid Cloud app instance using the Catalyst Cloud Portal when your Cisco ISE or pxGrid Cloud stops working or you need to update it.

This is an optional task. You should do this only if your Cisco ISE or pxGrid Cloud stops working or you need to update it.

Before you begin

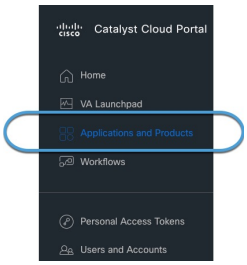
Make sure your current pxGrid Cloud identity source is active as discussed in Activate the pxGrid Cloud identity source, on page 2448.

Follow these steps to deactivate the pxGrid Cloud app instance:

Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

Step 2 In the Catalyst Cloud Portal, click ☰ > **Applications and Products** as the following figure shows:

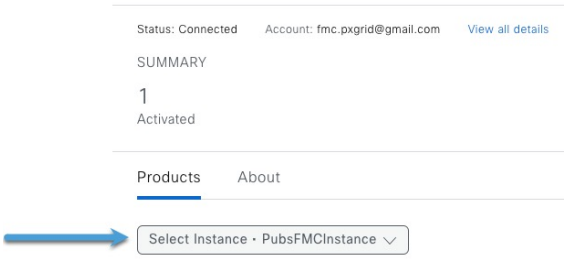


Step 3 Click **Applications**.

Step 4 Click **Manage** for Firewall Management Center.

From the **Select Instance** list, click the name of the firewall application you created earlier.

Example:



Step 5 In the Actions column, click More icon (···) > **Deactivate**.

Step 6 Wait until the product is removed.

You can click **Refresh** ( ) to see updated status if necessary.

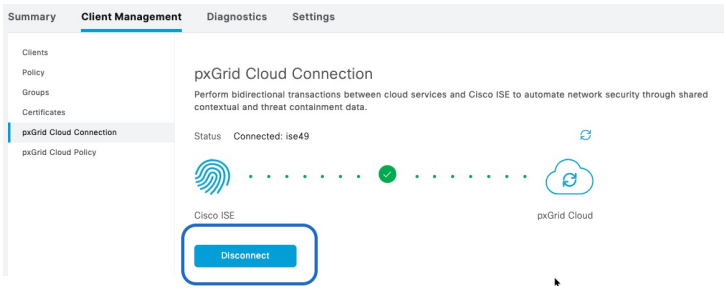
Step 7 ISE 3.3 or earlier: Disconnect the app instance:

a) Log in to Cisco ISE as an administrator.

b) Click **Administration** > **pxGrid Services** > **Client Management** > **pxGrid cloud connection**.

c) Click **Disconnect**.

Example:



Step 8 ISE 3.4 or later: Deregister the app instance:

a) Log in to Cisco ISE as an administrator.

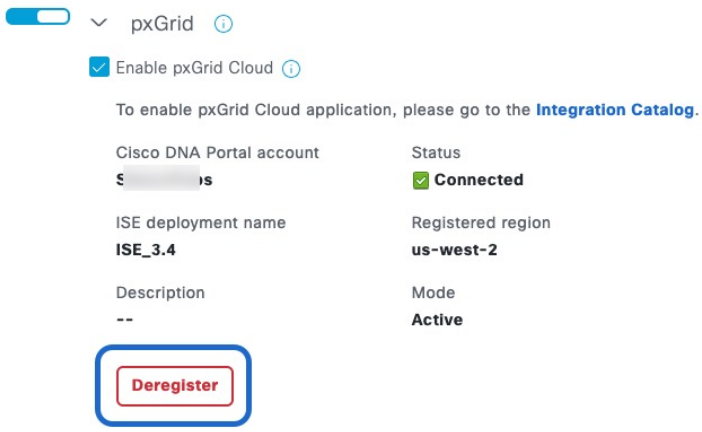
b) Click **Administration** > **System** > **Deployment**.

c) Expand **Deployment**.

d) Click the name of the ISE node.

e) In the **General Settings** tab page, scroll to locate **pxGrid**.
f) Click **Deregister**.

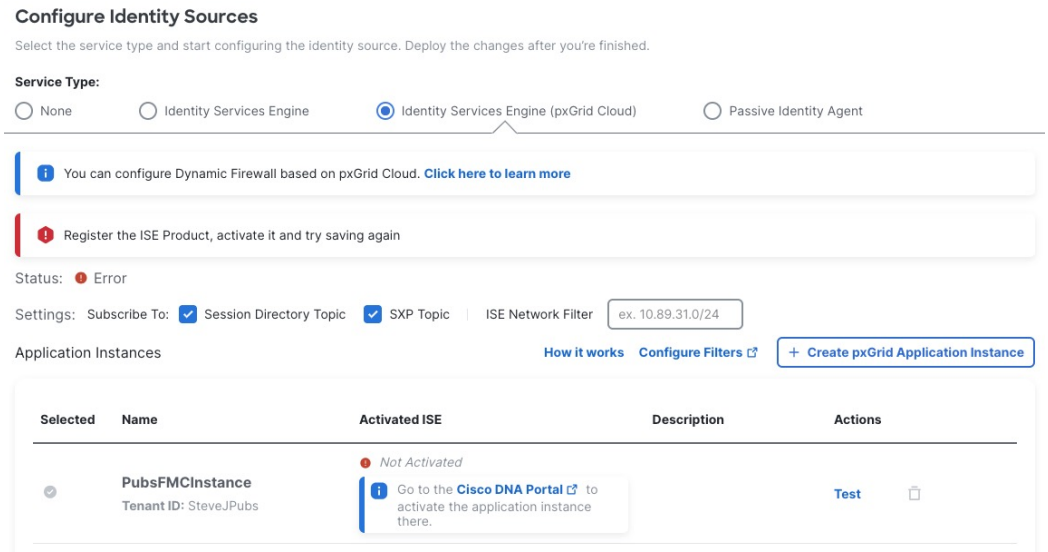
Example:



Step 9 To verify the app instance is deactivated in Security Cloud Control:

a) Log in to Security Cloud Control.
b) Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**
c) Click **Identity Services Engine (pxGrid Cloud)**.
d) Verify that **Not Activated** is displayed in the Activated ISE column.

Example:



What to do next

- To register Cisco ISE with pxGrid Cloud and activate the app instance, see:

- ISE 3.3 and earlier: Register Cisco ISE with the Catalyst Cloud Portal, on page 2439.
- ISE 3.4 and later: Create an app instance, on page 2473.

- To completely remove the identity source, see Delete the pxGrid Cloud identity source, on page 2462.

Delete the pxGrid Cloud identity source

Delete the pxGrid Cloud identity source from Security Cloud Control when you no longer need it.

This task explains how to delete the pxGrid Cloud identity source from Security Cloud Control, which is necessary if you do not want to use it again.

Before you begin

Deactivate the FMC app instance from the Catalyst Cloud Portal as discussed in Deactivate the pxGrid Cloud app instance, on page 2459.

Follow these steps to delete the pxGrid Cloud identity source:

Procedure

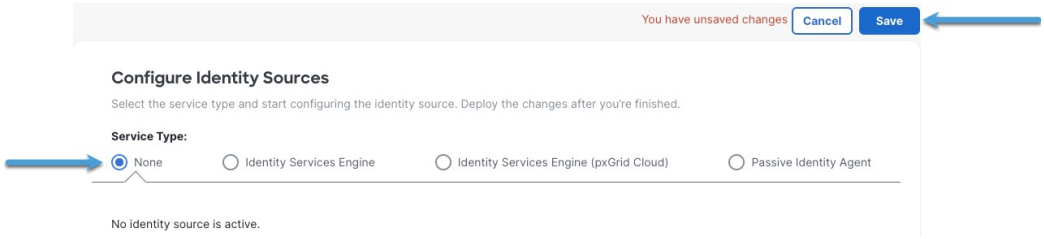
Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 2 Click **Identity Services Engine (pxGrid Cloud)**.

Step 3 For Service Type, click **None**.

Example:



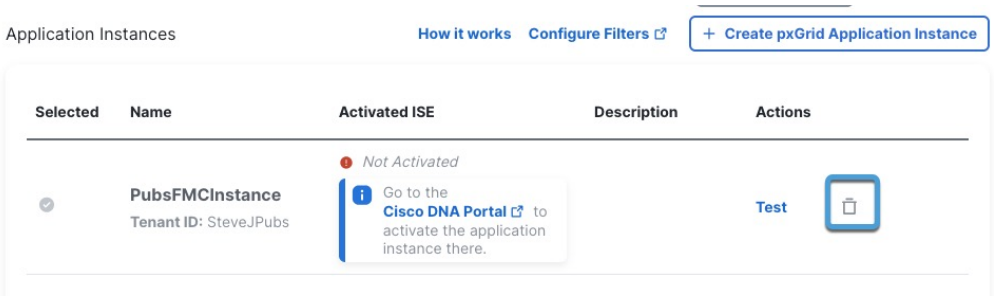
Step 4 Click **Save**.

Step 5 Confirm your choice when prompted.

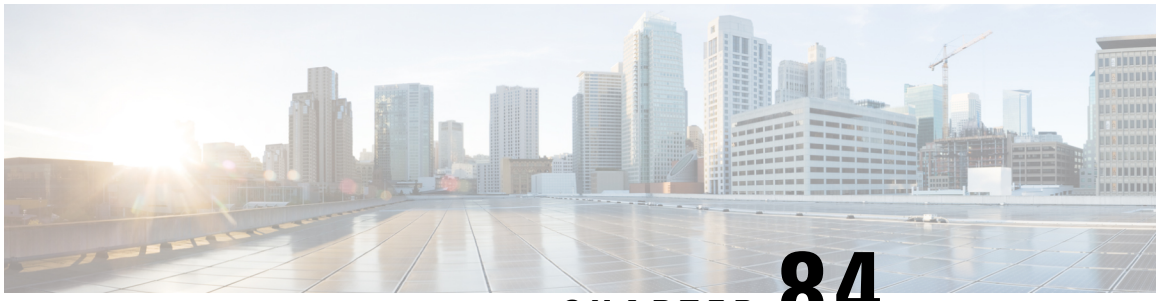
Step 6 Click **Identity Services Engine (pxGrid Cloud)**.

Step 7 Click **Delete** ( ).

Example:



Step 8 Confirm the action when prompted.



CHAPTER 84

Identity Source: pxGrid Cloud Identity (ISE 3.4 and Later)

The following topics discuss how to configure and use the pxGrid Cloud Identity Source with Cisco ISE version 3.4 and later.

- pxGrid Cloud identity source, on page 2465
- pxGrid Cloud identity source configuration, on page 2468
- Enable the pxGrid Cloud service in Cisco ISE, on page 2470
- Create an app instance, on page 2473
- Create the identity source, on page 2474
- Activate the app instance, on page 2475
- Activate the pxGrid Cloud identity source, on page 2478
- Test the pxGrid Cloud identity source, on page 2480
- Troubleshoot the pxGrid Cloud identity source, on page 2485
- Create dynamic attributes filters, on page 2486
- Create access control rules or DNS rules using dynamic attributes filters, on page 2487
- Deactivating and deleting the pxGrid Cloud identity source, on page 2489

pxGrid Cloud identity source

The Cisco Identity Services Engine (Cisco ISE) pxGrid Cloud identity source enables you to use subscription and user data from Cisco ISE in Cloud-Delivered Firewall Management Center access control rules. Also, the identity source uses constantly changing dynamic objects from Cisco ISE in access control policies in the Cloud-Delivered Firewall Management Center.

Integration technologies

The pxGrid Cloud identity source also uses:

- The Cisco Platform Exchange Grid (pxGrid), which enables multivendor, cross-platform network system collaboration in things like security monitoring and detection systems, network policy platforms, asset and configuration management, identity, and access management. pxGrid Cloud is the cloud-based interface to Cisco ISE.

  More information about pxGrid is available in resources such as What is PxGrid? on devnet.

- The Cisco Digital Network Architecture (Cisco DNA) delivers automation, security, predictive monitoring, and a policy-driven approach. It provides end-to-end network visibility and uses network insights to optimize network performance and deliver the best user and application experience.

  To use the pxGrid Cloud identity source with the Security Cloud Control, you must Create a Cisco Account.

  - What is pxGrid? on devnet
  - Cisco Platform Exchange Grid Cloud on devnet

Related Topics

Configuring a pxGrid Cloud identity source, on page 2468
Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436

Restrictions for the pxGrid Cloud identity source

Before you set up the pxGrid Cloud identity source, note these restrictions:

- pxGrid Cloud suppports these regions: `us-west-2`, `eu-central-1`, and `ap-southeast-1`.

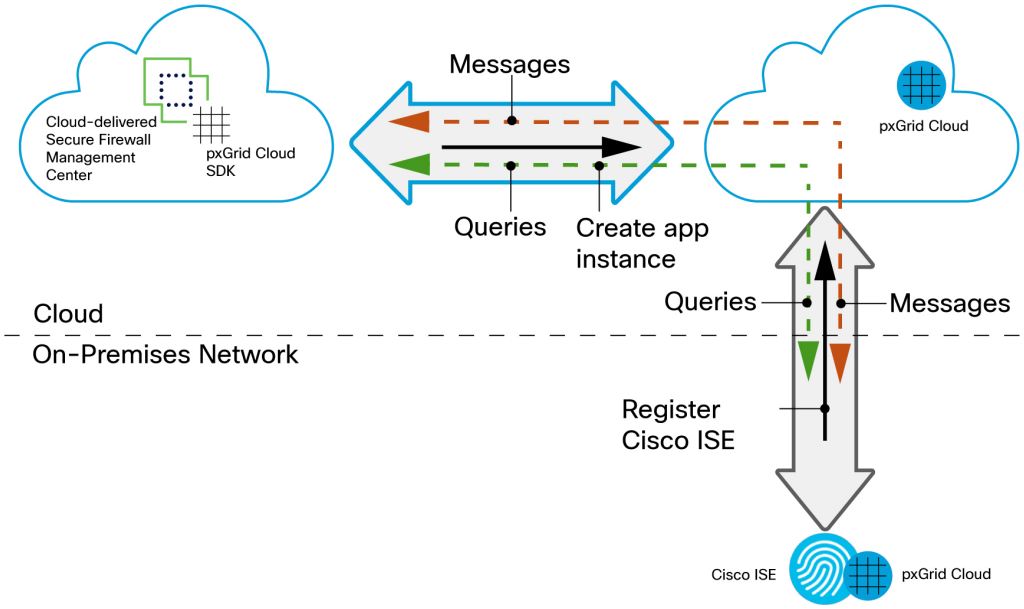
How the pxGrid Cloud identity source works

Summary

The key components involved in the pxGrid Cloud identity source are:

- Cloud-Delivered Firewall Management Center: Uses the pxGrid Cloud SDK to programmatically retrieve user information
- Cisco ISE: Provides user information, SGT, endpoint profile, and other details
- pxGrid Cloud: Facilitates secure data exchange between the management center and ISE
- Authentication process: Requires one-time passwords (OTP) to establish trusted communication

Workflow



These stages describe how the pxGrid Cloud identity source works:

1. In Cisco ISE, the administrator enables the use of pxGrid Cloud.
2. The administrator registers Cisco ISE as a product in pxGrid Cloud, which authenticates Cisco ISE and pxGrid Cloud and enables them to communicate with each other.

   The authentication process requires you to paste a one-time password (OTP) from pxGrid Cloud into Cisco ISE.
3. In pxGrid Cloud, the administrator creates an "app instance" that generates an OTP for use in the Cloud-Delivered Firewall Management Center to authenticate the two with each other.
4. After completing all the preceding tasks, the Cloud-Delivered Firewall Management Center (which includes the pxGrid Cloud SDK) can query Cisco ISE using pxGrid Cloud and retrieve sessions containing user information, SGT, endpoint profile, and other details.
5. Many types of dynamic objects can be filtered and sent to the Cloud-Delivered Firewall Management Center as dynamic objects to be used in access control rules. These include: SGT, endpoint profile, posture status, and machine authentication.

   User information is retrieved from Cisco ISE and group information is retrieved from either Microsoft Active Directory or Azure Active Directory.

Related Topics

Configuring a pxGrid Cloud identity source, on page 2468
Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436

pxGrid Cloud identity source configuration

These topics summarize how to configure a pxGrid Cloud identity source either for ISE 3.3 and earlier or for ISE 3.4 and later. The steps are different so make sure you follow them exactly.

Related Topics

Configuring a pxGrid Cloud identity source, on page 2468
Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436
Configuring a pxGrid Cloud identity source, on page 2468
Enable the pxGrid Cloud service in Cisco ISE, on page 2470
Create an app instance, on page 2473
Create the identity source, on page 2445
Activate the app instance, on page 2446
Activate the pxGrid Cloud identity source, on page 2448

Configuring a pxGrid Cloud identity source

Before you begin, create a Cisco Account.



Important *This topic applies to Cisco ISE version 3.4 or later*. If you are using an earlier version, see Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436 instead.

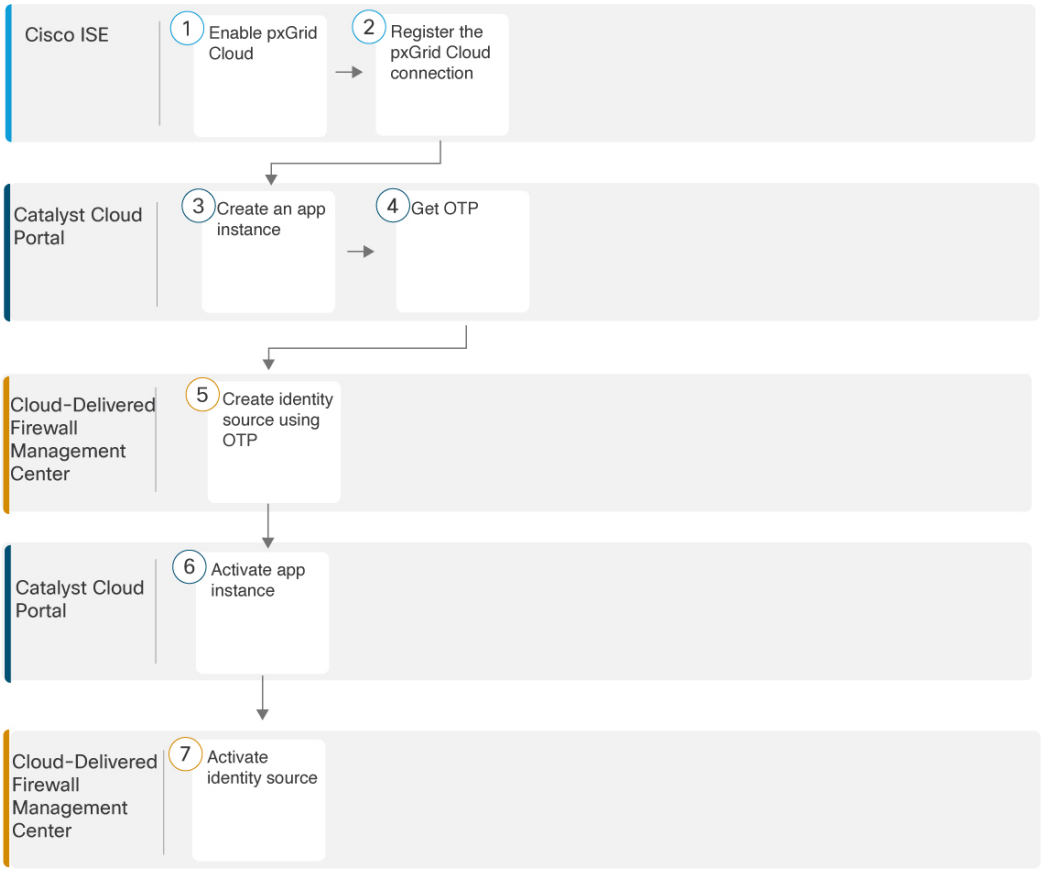
Summary

The key components involved in configuring a pxGrid Cloud identity source are:

- Cisco ISE: Enables pxGrid Cloud service and creates the connection
- Catalyst Cloud Portal: Creates and manages app instances for authentication
- Cloud-Delivered Firewall Management Center: Receives user data and creates identity sources for policy enforcement

Workflow

Figure 391: Configure a pxGrid Cloud identity source



1 Enable the pxGrid Cloud service in Cisco ISE, on page 2470
2 Register the pxGrid Cloud connection with Cisco ISE, on page 2442
3 Create an app instance, on page 2473
4 Create an app instance, on page 2473
5 Create a pxGrid Cloud identity source, on page 2443
6 Activate the app instance, on page 2446
7 Activate the pxGrid Cloud identity source, on page 2448

These are the stages of configuring a pxGrid Cloud identity source using Cisco ISE, the Catalyst Cloud Portal, and Cloud-Delivered Firewall Management Center:

1. Enable pxGrid Cloud in Cisco ISE.

   pxGrid Cloud enables you to subscribe to offers and to register apps (in this case, the Cloud-Delivered Firewall Management Center) for secure data exchange in a cloud environment.

   For more information, see Enable the pxGrid Cloud service in Cisco ISE, on page 2470.

2. Create an app instance and get the one-time password (OTP) required to create the pxGrid Cloud identity source in the Catalyst Cloud Portal.

   For more information, see Create an app instance, on page 2473.

3. Create the pxGrid Cloud identity source in the Cloud-Delivered Firewall Management Center.

   The identity source enables the Cloud-Delivered Firewall Management Center to authenticate with Cisco ISE and the Catalyst Cloud Portal so it can receive user data from Cisco ISE.

   For more information, see Create the identity source, on page 2445.

4. Activate the app instance in the Catalyst Cloud Portal.

   For more information, see Activate the app instance, on page 2446.

5. Activate the pxGrid Cloud identity source in the Cloud-Delivered Firewall Management Center.

   For more information, see Activate the pxGrid Cloud identity source, on page 2448.

What's next

After you complete this configuration process, you can:

- Test the pxGrid Cloud identity source to make sure it's working properly. For more information, see Test the pxGrid Cloud identity source, on page 2450.
- Create dynamic attributes filters, which define what dynamic objects are sent to the Cloud-Delivered Firewall Management Center. For more information, see Create dynamic attributes filters, on page 2455.
- Use the pxGrid Cloud identity source to enable dynamic objects, Microsoft AD user and groups, and Azure AD users and groups in access control rules.

Enable the pxGrid Cloud service in Cisco ISE

This task enables the pxGrid Cloud service in Cisco ISE, which allows your ISE deployment to connect to cloud-based pxGrid services for enhanced security and visibility.

This task explains how to enable pxGrid Cloud in Cisco ISE.

Before you begin

- Install and activate the Advantage license tier in your Cisco ISE deployment.
- The pxGrid Cloud agent creates an outbound HTTPS connection to Cisco pxGrid Cloud. Therefore, you must configure Cisco ISE proxy settings if the customer network uses a proxy to reach the internet. To configure proxy settings in Cisco ISE, click the **Menu** icon (☰) and choose **Administration > System > Settings > Proxy**.
- The Cisco ISE Trusted Certificates Store must include the root CA certificate required to validate the server certificate presented by Cisco pxGrid Cloud. Ensure that the **Trust for Authentication of Cisco Services** option is enabled for this root CA certificate. To enable this option, navigate to **Administration** > **System** > **Certificates**.

Follow these steps to enable the pxGrid Cloud service in Cisco ISE:

Procedure

Step 1 Log in as an administrator to Cisco ISE.

If your Cisco ISE server is part of a cluster, log in to the Primary Administration Node (PAN).

Step 2 Click the **Menu** icon (☰), then click **Administration** > **Deployment**.

Under Deployment Nodes, click the name of the deployment node.

Step 3 On the next page, scroll down to locate the **pxGrid** node.

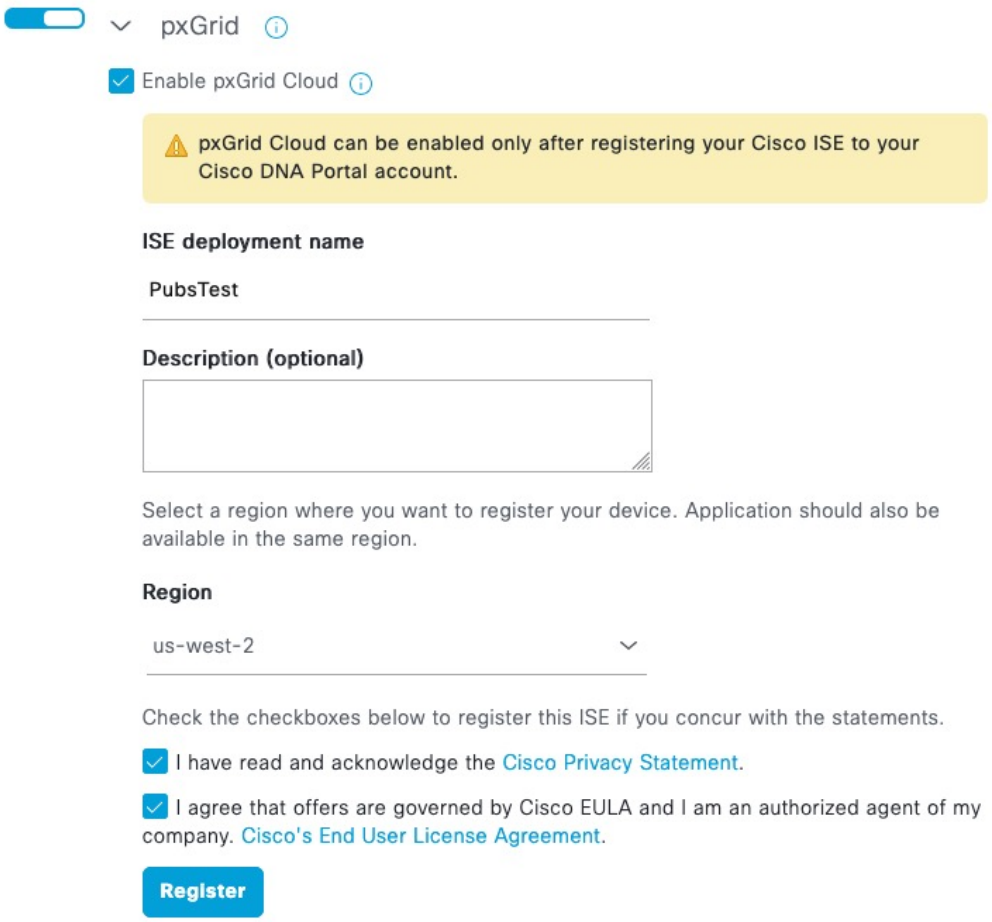
Select the **Enable pxGrid Cloud** check box.

Step 4 Enter the following information:

Option	Description
ISE Deployment Name	Enter a unique name to identify your ISE deployment.
Description (Optional)	Enter an optional description.
Region	Click the region you're in (**us-west-2**, **eu-central-1**, or **ap-southeast-1**). This is the only valid region at this time.
License agreements	You are required to select both check boxes to continue.

Step 5 Click **Register**.

Example:



Step 6 When prompted with an activation code for the Catalyst Cloud Portal, click **Next**.

Step 7 When prompted, log in to the Catalyst Cloud Portal.

Step 8 At the Select an Account dialog box, click the name of an account with which to register Cisco ISE.

Step 9 Click **Register ISE**.

What to do next

See Verify Cisco ISE registration in the Catalyst Cloud Portal, on page 2472.

Verify Cisco ISE registration in the Catalyst Cloud Portal

This task allows you to confirm that Cisco ISE is properly registered with the Catalyst Cloud Portal.

This task explains how to verify that Cisco ISE is registered with the Catalyst Cloud Portal.

Before you begin

Complete the tasks discussed in Enable the pxGrid Cloud service in Cisco ISE, on page 2470.

Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

Step 2 If prompted, select the account.

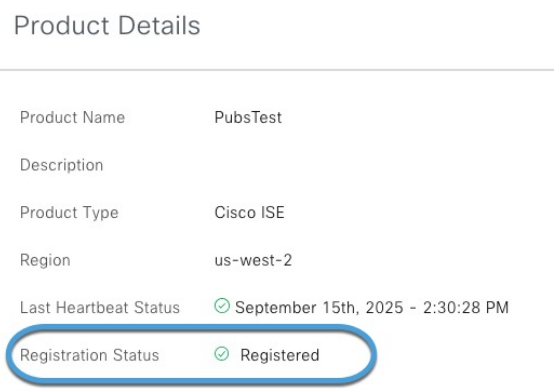
Step 3 In the left navigation, click ☰ > **Applications and Products**.

Step 4 At the top of the page, click **Products**.

Step 5 Click the name of the Cisco ISE deployment you created previously.

Step 6 Make sure that **Registered** is displayed for its status.

Example:



Create an app instance

Create an app instance to enable pxGrid Cloud identity source functionality for sending user session data to the Cloud-Delivered Firewall Management Center.

This task is one of several tasks you must perform to create a pxGrid Cloud identity source to send user session data to the Cloud-Delivered Firewall Management Center.

There is a one-hour time limit on the one-time password (OTP) required to complete this procedure successfully. You do not need to log in to Cisco ISE.

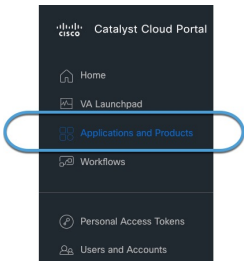
Before you begin

Verify Cisco ISE registration in the Catalyst Cloud Portal, on page 2472.

Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

Step 2 In the Catalyst Cloud Portal, click ☰ > **Applications and Products**.



Step 3 At the top of the page, click **Applications**.

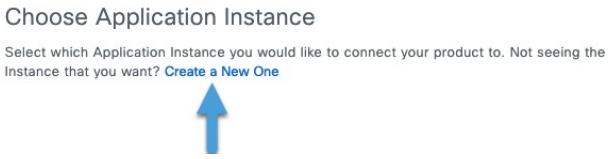
Step 4 From the **Regions** list, click **us-west-2**, **eu-central-1**, or **ap-southeast-1**.

Step 5 Click **Manage** (or **Activate**) next to **Firepower Management Center**.

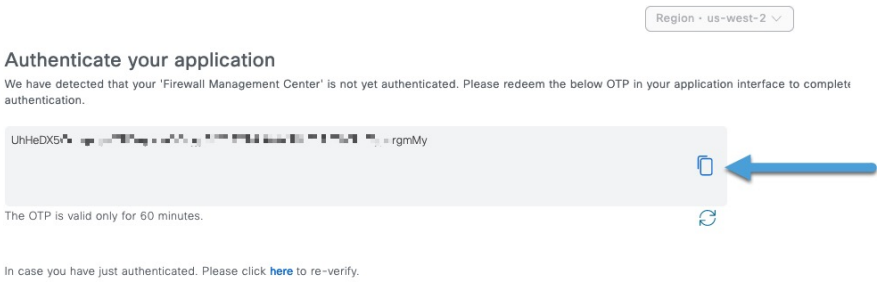
Step 6 Click **Add**.

Step 7 Click **Create a New One**.

The following figure shows an example.



Step 8 Click the copy button next to the displayed OTP.



Copy the OTP to a text file; it expires in 60 minutes.

Step 9 Continue with Create the identity source, on page 2445.

Create the identity source

This task creates a pxGrid Cloud identity source to send user session data to the Cloud-Delivered Firewall Management Center.

This task is one of several required to create a pxGrid Cloud identity source to send user session data to the Cloud-Delivered Firewall Management Center.

Before you begin

Complete the task discussed in Create an app instance, on page 2443.

Procedure

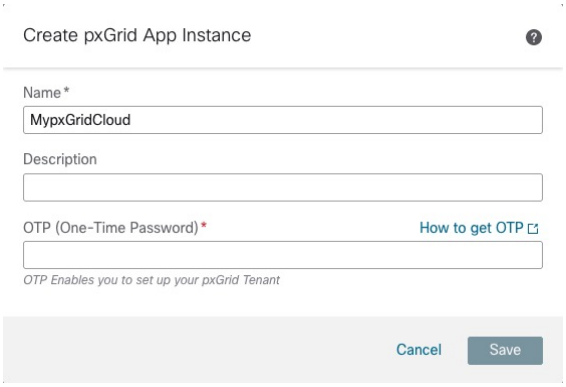
Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Step 2 Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**

Step 3 Click **Identity Services Engine (pxGrid Cloud)**.

Step 4 Click **Create pxGrid Application Instance**.

This figure shows an example.



Step 5 Enter the required information.

Value	Description
Name	Enter a name to uniquely identify this connector.
Description	Optional description.
OTP (One-Time Password)	Enter the OTP.

Step 6 Click **Create**.

Step 7 At the top of the page, click **Save**.

What to do next

Continue with Activate the app instance, on page 2446.

Activate the app instance

Activate an app instance to create a pxGrid Cloud identity source that sends user session data to the Cloud-Delivered Firewall Management Center.

There is a one-hour time limit on the one-time password (OTP) required to complete this procedure successfully. You do not need to log in to Cisco ISE.

Before you begin

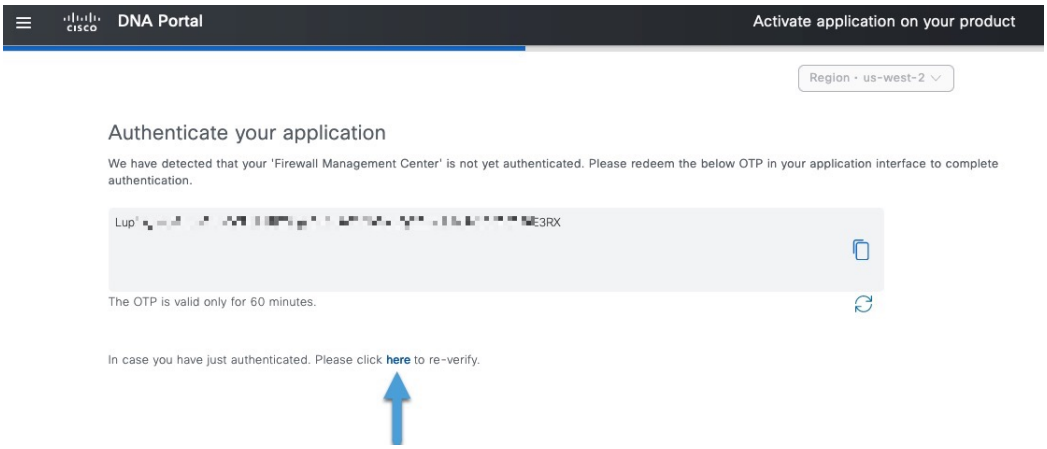
Complete the task discussed in Create the identity source, on page 2445.

Follow these steps to activate the app instance:

Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

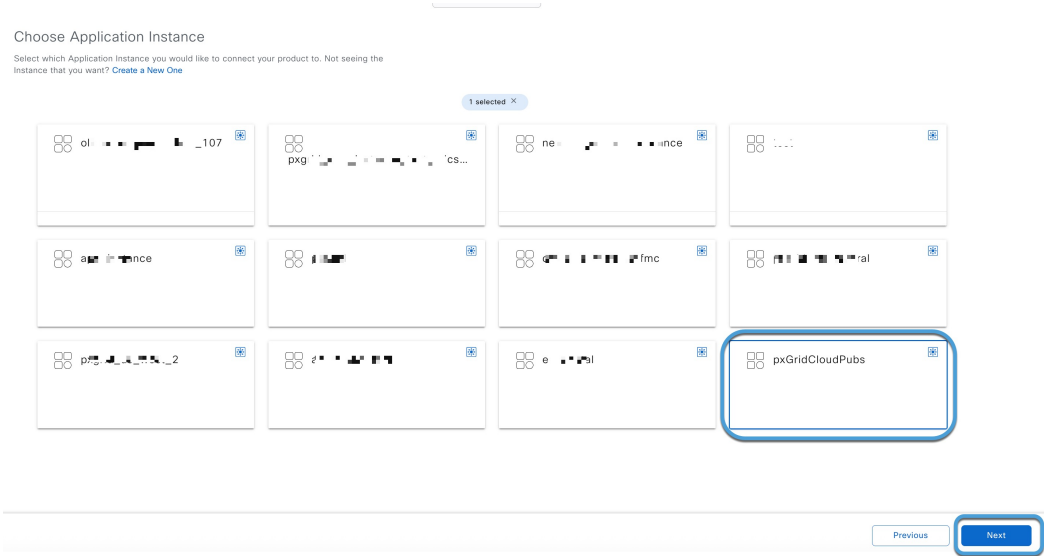
Step 2 Reverify the app by clicking the word **here** as the following figure shows.



Step 3 Click the name of the application instance you just created in Cloud-Delivered Firewall Management Center.

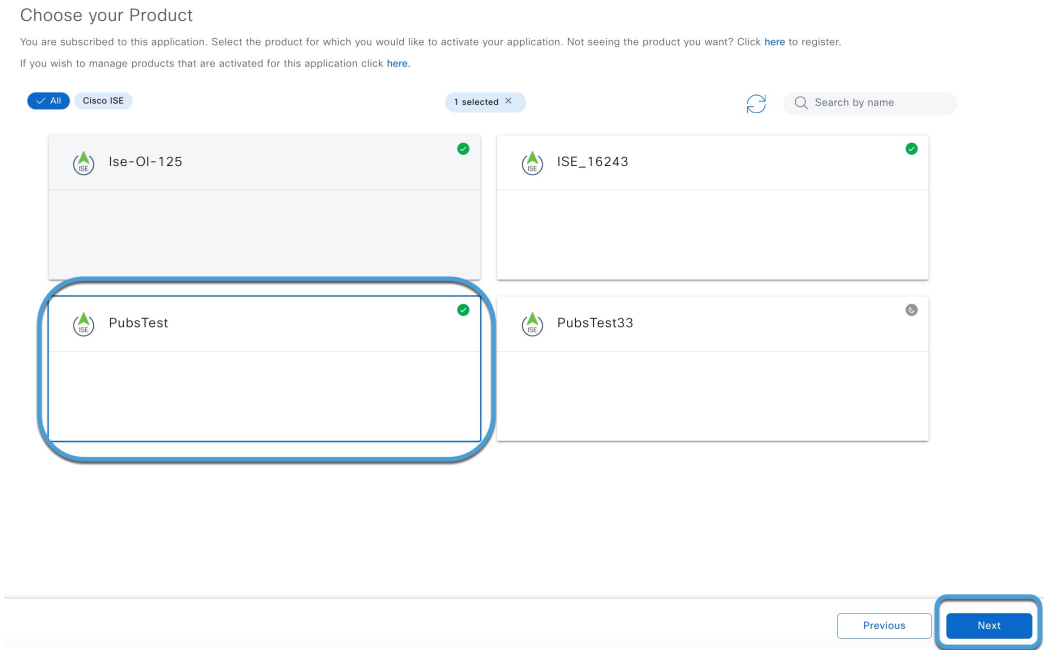
Step 4 Click **Next**.

Example:

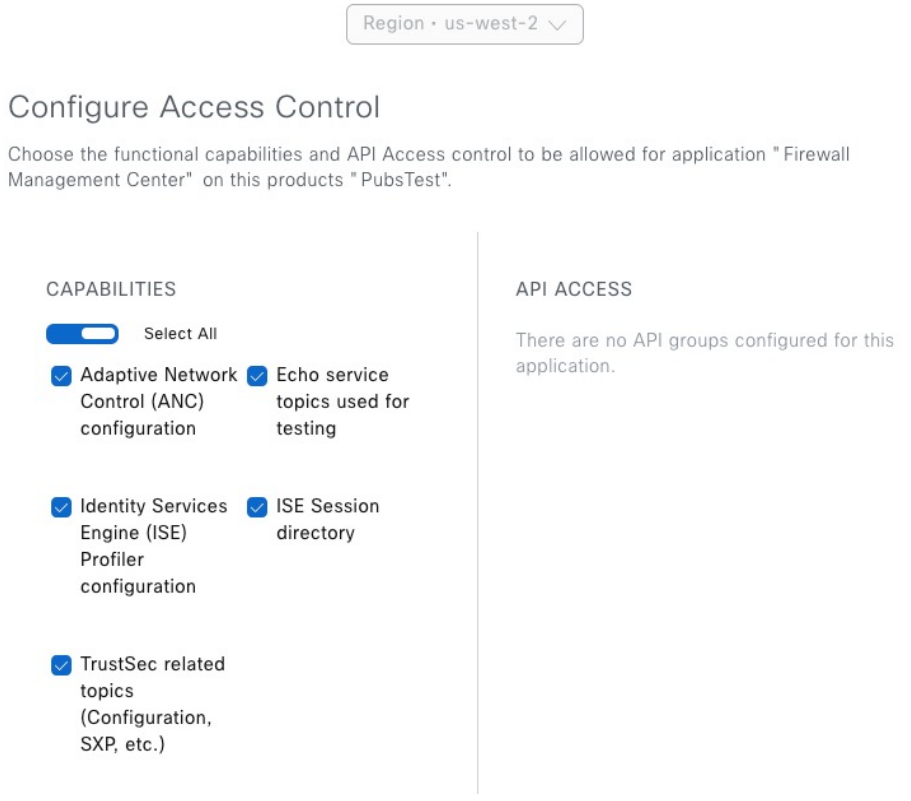


Step 5 On the Choose Product page, click the name of the Cisco ISE product and click **Next**.

Example:



Step 6 Select the check box next to each scope.
The following figure shows an example.



Step 7 Click **Next**.

Step 8 Review the displayed information for accuracy. Make sure all scopes are selected.

Step 9 Click **Activate**.

It can take several minutes for the app instance to be activated.

Continue with Activate the pxGrid Cloud identity source, on page 2448.

Activate the pxGrid Cloud identity source

This task activates the pxGrid Cloud identity source to enable identity services integration in your security infrastructure.

This task explains how to activate the pxGrid Cloud identity source in the Security Cloud Control.

Before you begin

Complete the tasks discussed in Activate the app instance, on page 2446.

Note Only one pxGrid Cloud identity source can be active at a time.

Procedure

Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

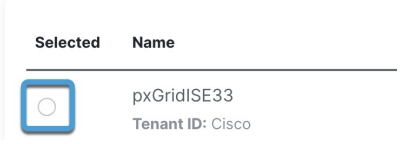
Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 2 Click **Identity Services Engine (pxGrid Cloud)**.

Step 3 Click **Save** at the top of the page.

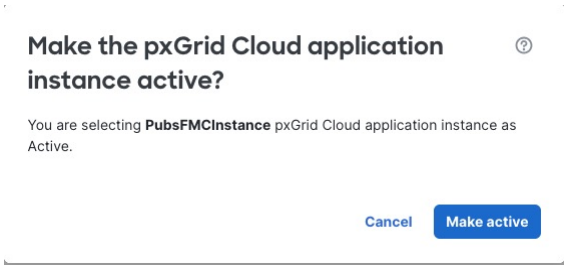
Step 4 If a green check mark is *not* displayed next to the name of the identity source, select it.

Example:



Step 5 Click **Make Active**.

Example:



Step 6 (Optional.) Select the following options if desired:

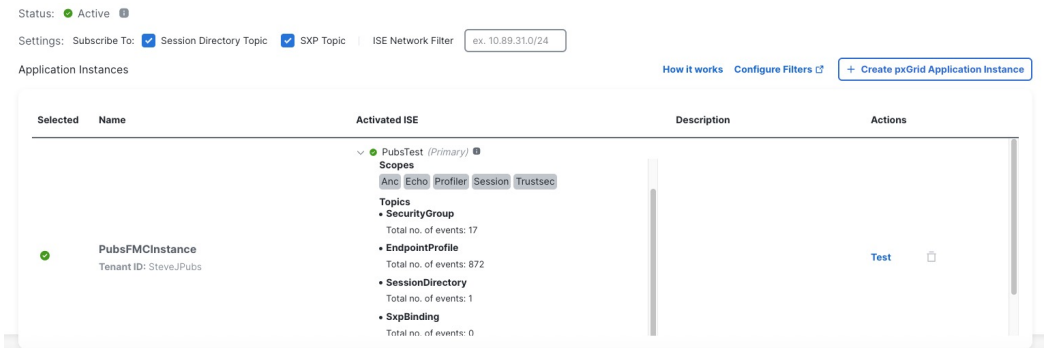
- **Session Directory Topic**: Select the check box to receive ISE user session information from the Cisco ISE server.
- **SXP Topic**: Select the check box to receive updates to SGT-to-IP mappings when available from the ISE server. This option is required to use destination SGT tagging in access control rules.
- **ISE Network Filter**: Optional filter you can set to restrict the data that Cisco ISE reports. If you provide a network filter, Cisco ISE reports data from the networks in that filter.

  You have the following options:

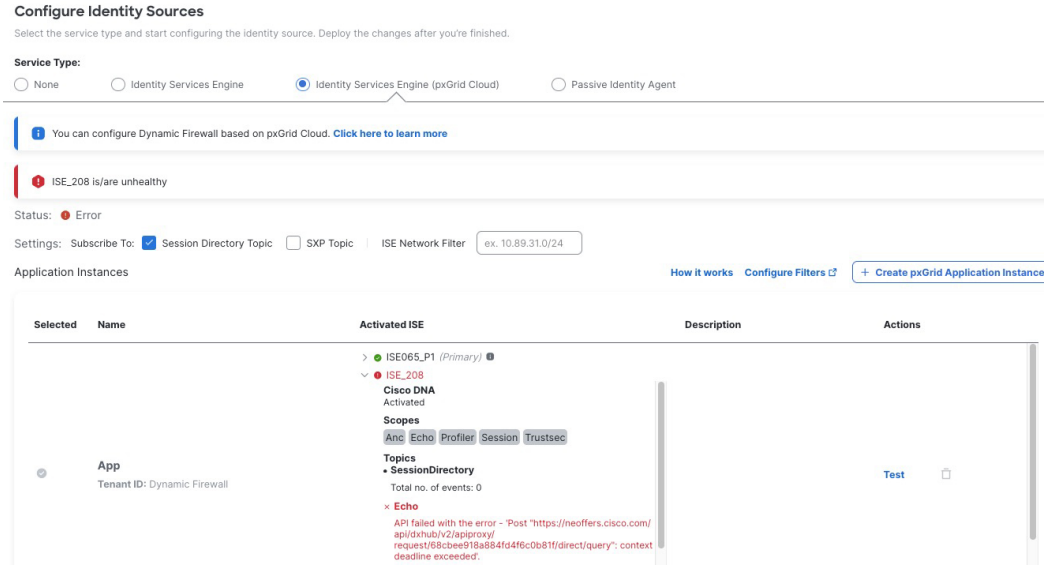
  - Leave the field blank to specify **`any`**.
  - Enter a single IPv4 address block using CIDR notation.
  - Enter a list of IPv4 address blocks using CIDR notation, separated by commas.

Step 7 Under Activated ISE, expand the identity source.

Example *normal* result:



Example *error* result:



In the event of an error, see Test the pxGrid Cloud identity source, on page 2450.

Step 8 Verify the status is Active and that all scopes and topics are displayed.

Wait a few minutes for data to be downloaded.

What to do next

See Test the pxGrid Cloud identity source, on page 2450.

Test the pxGrid Cloud identity source

This task enables you to perform diagnostics using the Security Cloud Control to determine if the identity source is working.

Errors might include communication with Cisco ISE, or with the Cisco ISE configuration with Catalyst Cloud Portal.

Procedure

Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Step 2

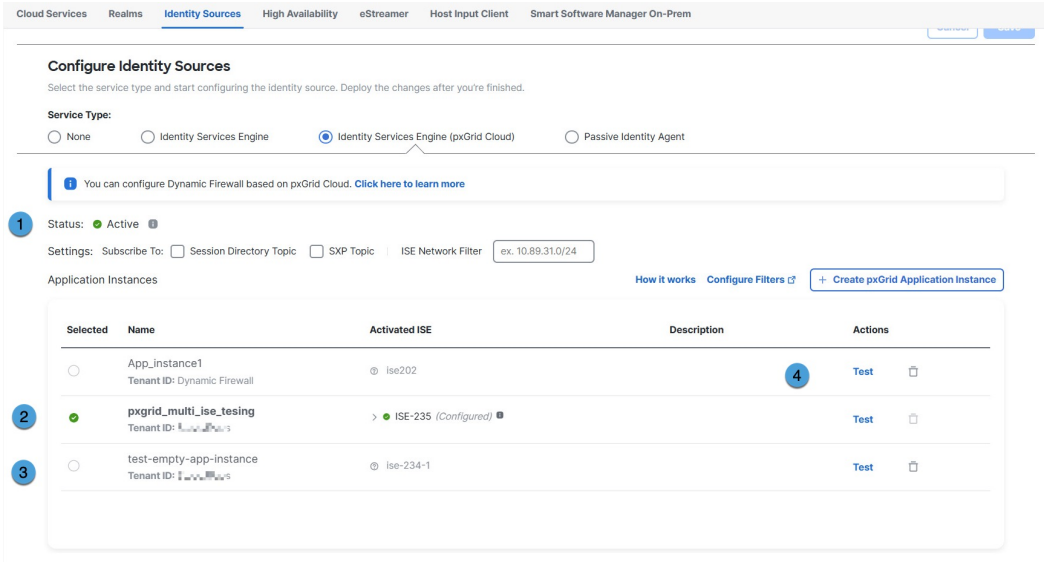
Step 3 Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 4

Step 5 Click **Identity Services Engine (pxGrid Cloud)**.

Step 6 Review the configuration status information.

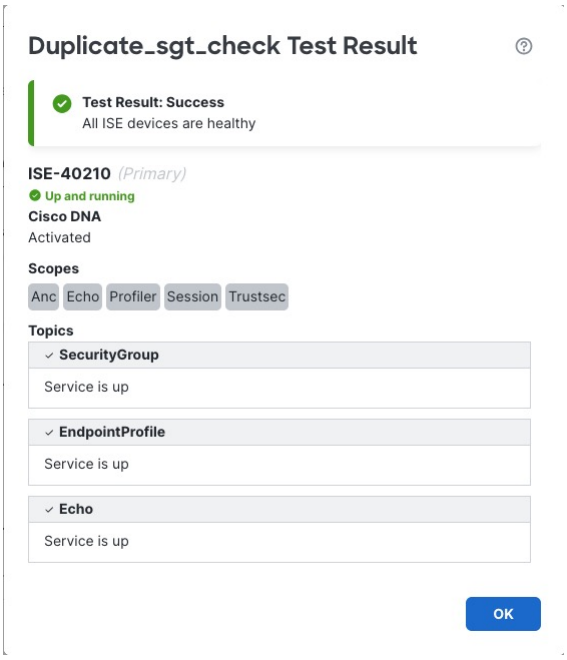
The following figure shows an example configuration.



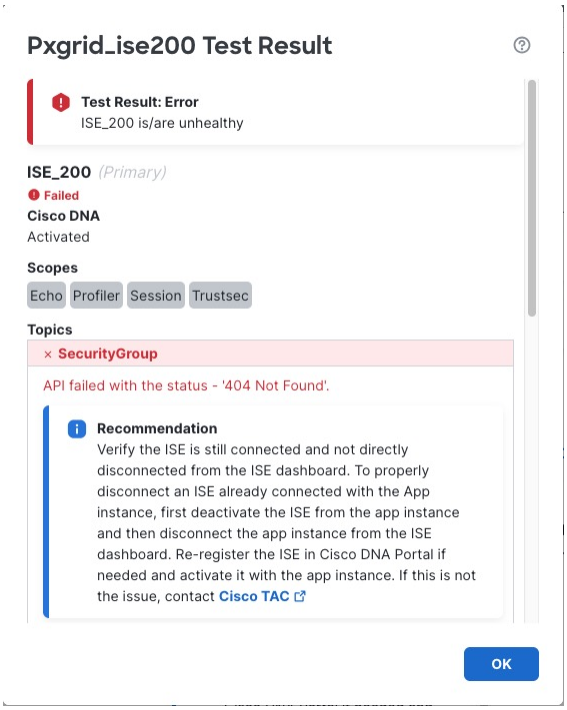
The following table has more information about the numbered areas in the figure.

Number	Meaning
1	Overall status Any errors in the overall status of the Cisco ISE app instances are displayed. In that case, scroll to that instance and either expand the error message or click **Test** for more information.
2	Active A green check mark indicates the app is active.
3	Inactive A dimmed app instance is inactive. You can activate it by selecting the check box next to its name and then clicking **Make active**.
4	**Test** button Click **Test** to perform diagnostic tests that show more detailed status of the app instance.

The following figure shows a sample success message.



The following figure shows an example error result.



Step 7 Click **Test** to perform diagnostic tests.

If errors occur, use the following error code reference to help diagnose and solve issues with Cisco ISE, pxGrid Cloud, and the Catalyst Cloud Portal. If these suggestions do not work, or if you have a different issue, contact Cisco TAC.

Table 133: Error code troubleshooting

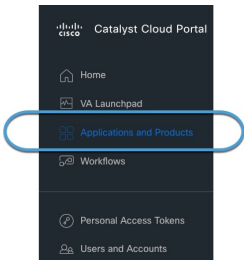
Error Code	Troubleshooting Steps
403 – Forbidden	Verify the Cisco ISE product is not in a **Pending** or **Suspended** state in the Catalyst Cloud Portal. If suspended, verify that Cisco ISE is registered as discussed in Enable pxGrid Cloud service in Cisco ISE and register your device. Additionally, verify pxGrid Cloud services are publicly available.
404 – Not Found	Verify the Cisco ISE server is not directly disconnected from the Cisco ISE dashboard. To properly disconnect Cisco ISE already connected with the app instance, first deactivate Cisco ISE from the app instance and then disconnect the app instance from the Cisco ISE dashboard.
408 – Request Timeout	Check whether there are any general connectivity issues with Cisco ISE and verify pxGrid Cloud connectivity status is **Connected** in the ISE dashboard under ☰ > **Administration** > **pxGrid Services** > **Client Management** > **pxGrid Cloud Connection**. Verify the Cisco ISE server is not directly disconnected from the Cisco ISE dashboard.
413 – Content Too Large	Review the pxGrid Cloud API limitations on GitHub. If needed, consider upgrading your Cisco ISE version to fully utilize pxGrid Cloud support.
500 – Internal Server Error	Check that the Cisco ISE server is operational and that pxGrid Cloud services are active (verify MNT, SXP, pxGrid nodes, and so on). For more information, see Monitoring and debugging in the Cisco pxGrid chapter in the *Cisco Identity Services Engine Administrator Guide*.

Step 8 If the product is in a **Pending** or **Suspended** state, verify that it is active.

a) Log in to the Catalyst Cloud Portal.

b) In the Catalyst Cloud Portal, go to ☰ > **Applications and Products**.

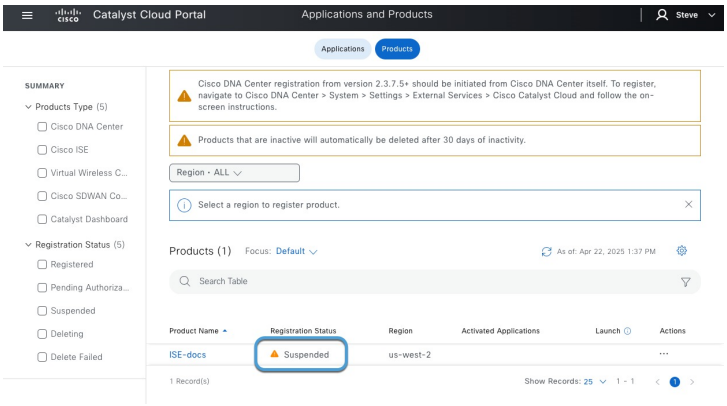
Example:



c) Click the **Products** tab.

Example:

The following figure shows an example of a suspended product.

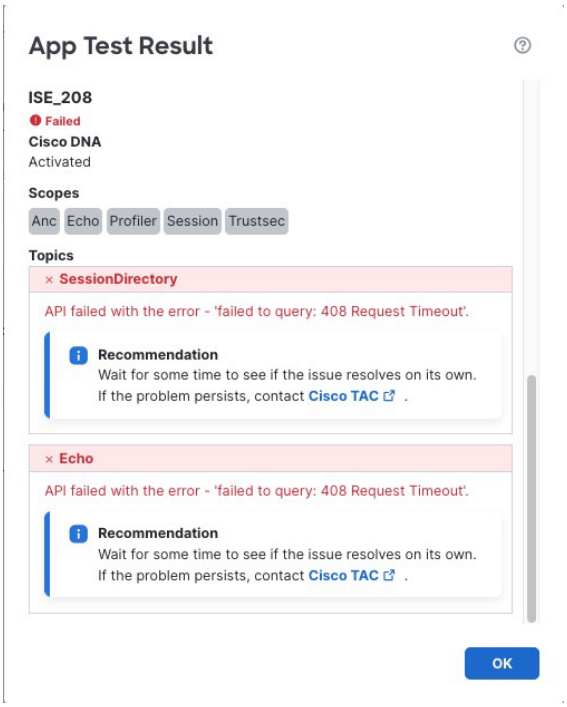


d) To correct the issue, in the Actions column, click ••• and click **Generate OTP**.

e) Use the OTP as discussed in Create the identity source, on page 2445.

Step 9 If you encounter a cluster member not reachable error, find what node is not reachable.

If a member of the Cisco ISE cluster is not reachable, a page like this is displayed:



To find what node is not reachable, log in to Cisco ISE primary administration node as an administrator and click click the **Menu** icon (☰) and choose **Administration** > **System** > **Deployment**, then see Node Status in a Cisco ISE Deployment.

Troubleshoot the pxGrid Cloud identity source

These topics describe troubleshooting the pxGrid Cloud identity source.

Related Topics

Primary device cannot be processed, on page 2458
Primary device cannot be processed, on page 2458

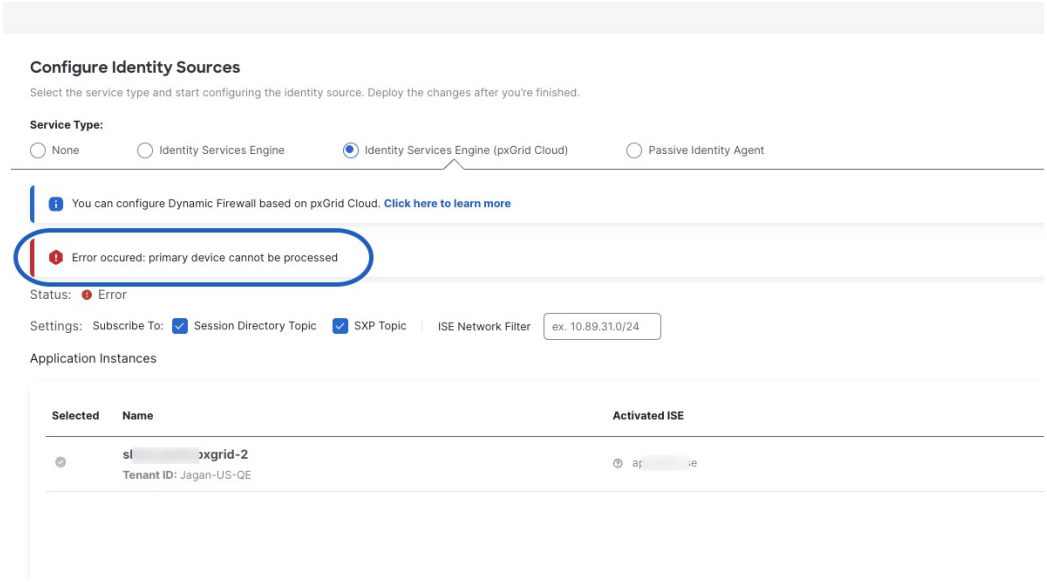
Primary device cannot be processed

This task resolves the error "primary device cannot be processed" or "ISE is unhealthy" that occurs when a Cisco ISE cluster is associated with more than one app instance.

Each Cisco ISE cluster must be associated with one and only one app instance, typically in a single dedicated tenant.

If you associate a Cisco ISE with more than one app instance, an error such as `Error occured: primary device cannot be processed` or `ISE is unhealthy` is displayed for the identity source.

Example:



The solution is to deactivate the ISE properly from other app instances for that tenant, before using it in any other tenant or app instance

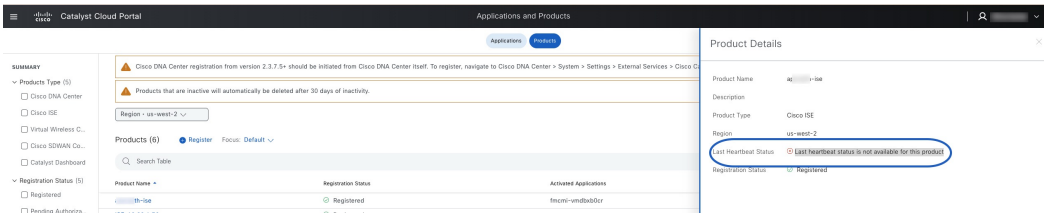
Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

Step 2 At the top of the page, click **Applications**.

Step 3 Locate a Cisco ISE product that is activated and verify it is the one causing the issue.

Example:



Step 4 Wait for the product to be removed.

Step 5 Deactivate the app instance as described in Deactivate the pxGrid Cloud app instance, on page 2459.

Create dynamic attributes filters

Create dynamic attributes filters to expose them as dynamic objects in Security Cloud Control for use in access control policies.

Dynamic attributes filters that you define using the Dynamic Attributes Connector are exposed in the Security Cloud Control as dynamic objects that can be used in access control policies. For example, restrict access to an AWS server for the Finance Department to only members of the Finance group defined in Microsoft Active Directory.

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

For more information about access control rules, see Create access control rules or DNS rules using dynamic attributes filters, on page 2456.

Before you begin

Connectors, on page 2631

Procedure

Step 1 Log in to Security Cloud Control.

a) Click **Firewall**.

b) Click **Administration** > **Dynamic Attributes Connector** > **Dynamic Attributes Filters**.

Step 2 Do any of the following:

- Add a new filter: click **Add** ( ).
- Edit or delete a filter: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 3 Enter the following information.

Item	Description
Name	Unique name to identify the dynamic filter (as a dynamic object) in a policy and in the Security Cloud Control Object Manager (**External Attributes** > **Dynamic Object**).
Connector	From the list, click the name of a connector to use.
Query	Click Add .

Step 4 To add or edit a query, enter the following information.

Item	Description
Key	Click a key from the list. Keys are fetched from the connector.
Operation	Click one of the following: • **Equals** to exactly match the key to the value. • **Contains** to match the key to the value if any part of the value matches.
Values	Click either **Any** or **All** and click one or more values from the list. Click **Add another value** to add values to your query.

Step 5 Click **Show Preview** to display a list of networks or IP addresses returned by your query.

Step 6 When you're finished, click **Save**.

Step 7 (Optional.) Verify the dynamic object in the Security Cloud Control.

a) Log in to the Security Cloud Control.

b) Click **Policies** > **Firewall Threat Defense**.

c) Click **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.

The dynamic attribute query you created should be displayed as a dynamic object.

Create access control rules or DNS rules using dynamic attributes filters

Create access control rules using dynamic objects that are based on dynamic attributes filters you have previously configured.

This topic discusses how to create access control rules using dynamic objects (these dynamic objects are named after the dynamic attributes filters you created previously).

To add dynamic attributes filters to DNS policies, see Create basic DNS policies, on page 2561.

To add dynamic attributes filters to DNS policies, see Creating Basic DNS Policies.

Before you begin

Create dynamic attributes filters as discussed in Create dynamic attributes filters, on page 2455.

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

Follow these steps to create access control rules or DNS rules using dynamic attributes filters:

Procedure

Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Click **Policies** > **Firewall Threat Defense** > **Access Control heading** > **Access Control**.

Step 2 Click **Firewall**.

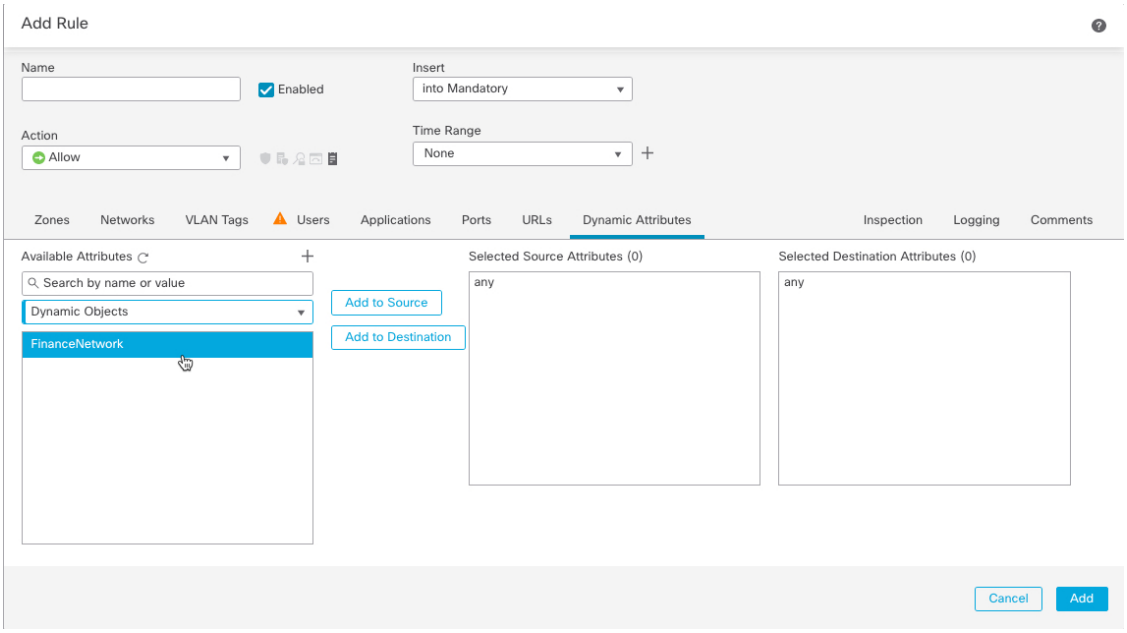
Step 3 Click **Edit** ( ) next to an access control policy.

Step 4 Click **Add Rule**.

Step 5 Click the **Dynamic Attributes** tab.

Step 6 In the Available Attributes section, from the list, click **Dynamic Objects**.

The following figure shows an example.



This example shows a dynamic object named `APIC Dynamic Attribute` that corresponds to the dynamic attribute filter created in the dynamic attributes connector.

Step 7 Add the desired object to source or destination attributes.

Step 8 Add other conditions to the rule if desired.

See Dynamic attributes rule conditions, on page 2681.

Deactivating and deleting the pxGrid Cloud identity source

Deactivating the FMC app instance in the Catalyst Cloud Portal is an optional task to troubleshoot issues with the Cisco ISE integration.

You should delete the pxGrid Cloud identity source from the Security Cloud Control when you are certain you don't want to use it again.

Related Topics

Deactivate the pxGrid Cloud app instance, on page 2459
Delete the pxGrid Cloud identity source, on page 2462

Deactivate the pxGrid Cloud app instance

This task explains how to deactivate a pxGrid Cloud app instance using the Catalyst Cloud Portal when your Cisco ISE or pxGrid Cloud stops working or you need to update it.

This is an optional task. You should do this only if your Cisco ISE or pxGrid Cloud stops working or you need to update it.

Before you begin

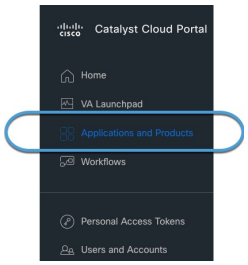
Make sure your current pxGrid Cloud identity source is active as discussed in Activate the pxGrid Cloud identity source, on page 2448.

Follow these steps to deactivate the pxGrid Cloud app instance:

Procedure

Step 1 Log in to the Cisco Catalyst Cloud Portal.

Step 2 In the Catalyst Cloud Portal, click ☰ > **Applications and Products** as the following figure shows:

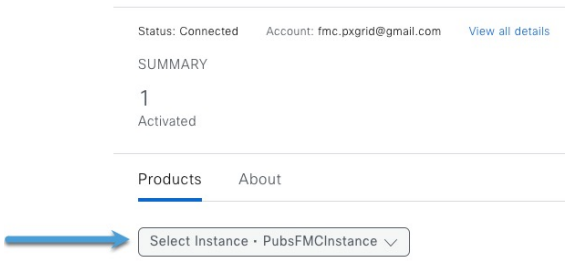


Step 3 Click **Applications**.

Step 4 Click **Manage** for Firewall Management Center.

From the **Select Instance** list, click the name of the firewall application you created earlier.

Example:



Step 5 In the Actions column, click More icon (···) > **Deactivate**.

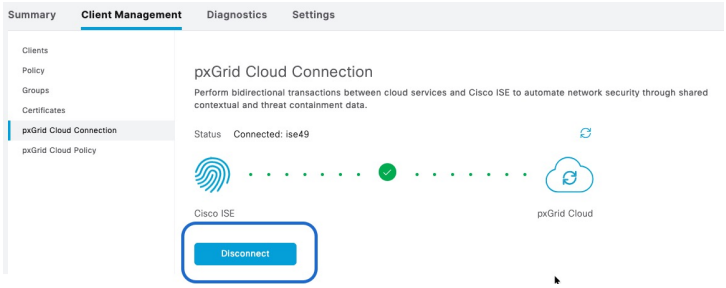
Step 6 Wait until the product is removed.

You can click **Refresh** ( ) to see updated status if necessary.

Step 7 ISE 3.3 or earlier: Disconnect the app instance:

a) Log in to Cisco ISE as an administrator.
b) Click **Administration** > **pxGrid Services** > **Client Management** > **pxGrid cloud connection**.
c) Click **Disconnect**.

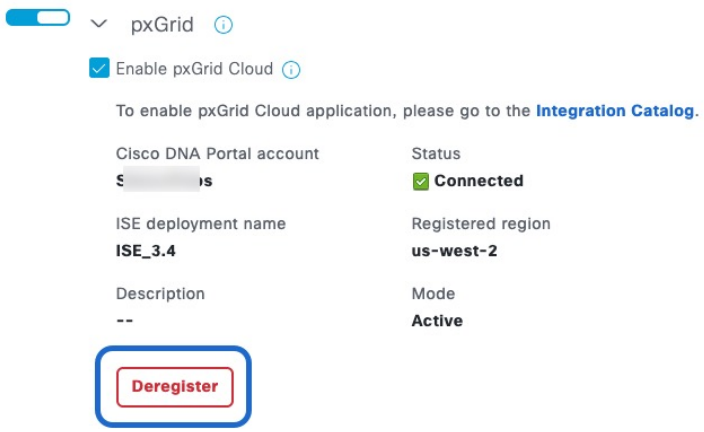
Example:



Step 8 ISE 3.4 or later: Deregister the app instance:

a) Log in to Cisco ISE as an administrator.
b) Click **Administration** > **System** > **Deployment**.
c) Expand **Deployment**.
d) Click the name of the ISE node.
e) In the **General Settings** tab page, scroll to locate **pxGrid**.
f) Click **Deregister**.

Example:



Step 9 To verify the app instance is deactivated in Security Cloud Control:

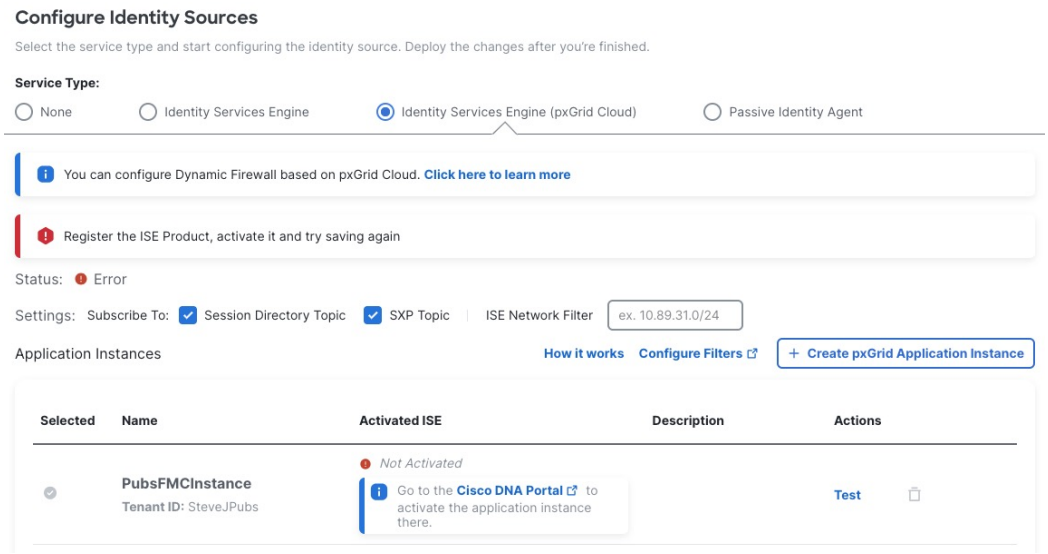
a) Log in to Security Cloud Control.

b) Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**

c) Click **Identity Services Engine (pxGrid Cloud)**.

d) Verify that **Not Activated** is displayed in the Activated ISE column.

Example:



What to do next

- To register Cisco ISE with pxGrid Cloud and activate the app instance, see:
  - ISE 3.3 and earlier: Register Cisco ISE with the Catalyst Cloud Portal, on page 2439.
  - ISE 3.4 and later: Create an app instance, on page 2473.
- To completely remove the identity source, see Delete the pxGrid Cloud identity source, on page 2462.

Delete the pxGrid Cloud identity source

Delete the pxGrid Cloud identity source from Security Cloud Control when you no longer need it.

This task explains how to delete the pxGrid Cloud identity source from Security Cloud Control, which is necessary if you do not want to use it again.

Before you begin

Deactivate the FMC app instance from the Catalyst Cloud Portal as discussed in Deactivate the pxGrid Cloud app instance, on page 2459.

Follow these steps to delete the pxGrid Cloud identity source:

Procedure

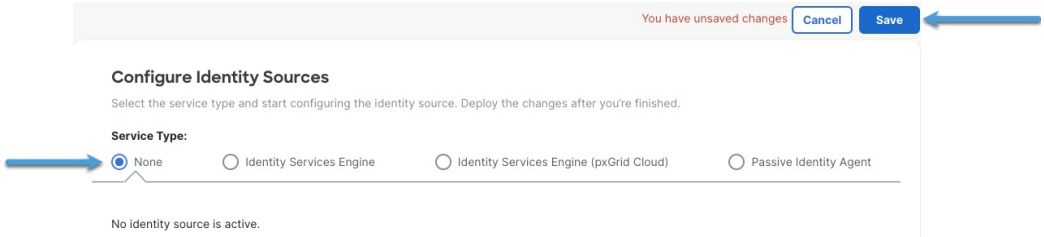
Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Click **Policies** > **Threat Defense** > **Integration** > **Other Integrations** > **Identity** > **Identity Sources**.

Step 2 Click **Identity Services Engine (pxGrid Cloud)**.

Step 3 For Service Type, click **None**.

Example:



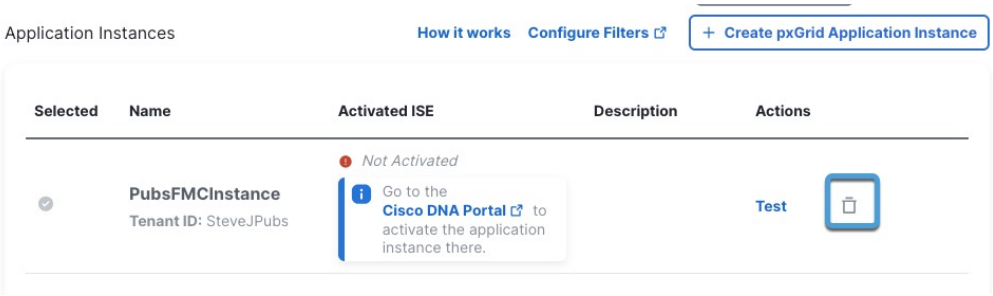
Step 4 Click **Save**.

Step 5 Confirm your choice when prompted.

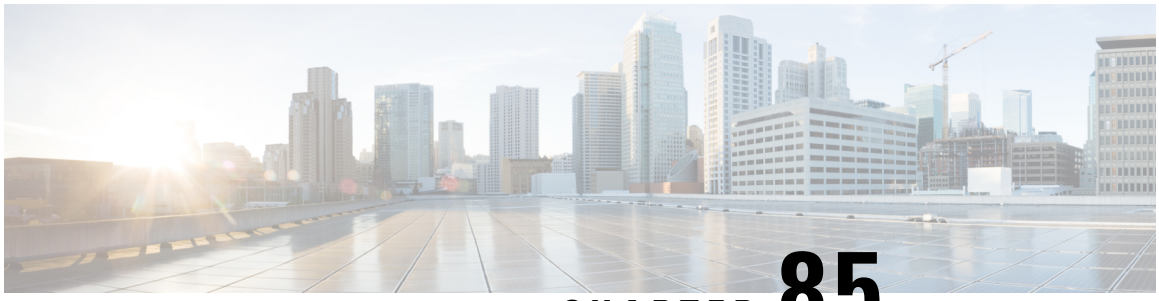
Step 6 Click **Identity Services Engine (pxGrid Cloud)**.

Step 7 Click **Delete** ( ).

Example:



Step 8 Confirm the action when prompted.



CHAPTER 85

Identity Source: Remote Access VPN

The following topics discuss how to perform user awareness and user control with Remote Access VPN:

- Remote access VPN identity sources, on page 2495
- Configure Remote Access VPN for AD-authenticated users, on page 2496
- Hybrid Entra ID and on-premises AD deployments, on page 2497
- Troubleshooting guidelines for remote access VPN identity source, on page 2497

Remote access VPN identity sources

A remote access VPN identity source authenticates a user before establishing secure remote connection to the network. It requires the user to provide credentials through a client to verify the identity.

Identity source configuration

Secure Client is the only client supported on endpoint devices for remote VPN connectivity to Firewall Threat Defense devices.

When you set up a secure VPN gateway as discussed in Configure a remote access VPN policy, you can set up an identity policy for those users and associate the identity policy with an access control policy, provided your users are in an Active Directory repository.

Note If you use remote access VPN with User Identity and RADIUS as the identity source, you must configure the realm (**Objects** > **Object Management** > **AAA Server** > **RADIUS Server Group**).

The login information provided by a remote user is validated by an LDAP or AD realm or a RADIUS server group. These entities are integrated with the Secure Firewall Threat Defense secure gateway.

Note If users authenticate with remote access VPN using Active Directory as the authentication source, users must log in using their username; the format domain\username or username@domain fails. (Active Directory refers to this username as the logon name or sometimes as sAMAccountName.) For more information, see User Naming Attributes on MSDN.

If you use RADIUS to authenticate, users can log in with any of the preceding formats.

Once authenticated through a VPN connection, the remote user takes on a VPN identity. This VPN identity is used by identity policies on the Secure Firewall Threat Defense secure gateway to recognize and filter network traffic belonging to that remote user.

Identity policies are associated with access control policies, which determine who has access to network resources. It is in this way that a remote user blocked or allowed to access your network resources.

Related Topics

Secure Connections Overview, on page 1709
Remote access VPN
VPN basics, on page 1710
Configure a remote access VPN policy

Configure Remote Access VPN for AD-authenticated users

Configure remote access VPN for Active Directory-authenticated users. Manage VPN connections and apply identity-based policies to control user access and traffic filtering.

Remote access VPN user control allows you to implement identity-based policies that manage and monitor VPN user connections, providing enhanced security and access control for your network environment.

Before you begin

- Create a realm as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
- To use authentication, authorization, and auditing (AAA), set up a RADIUS server group as discussed in Add a RADIUS server group, on page 2943.

Follow these steps to configure RA VPN for user control:

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Devices** > **VPN** > **Remote Access**.

Step 3 See Configure a remote access VPN policy.

Step 4 Specify users to control and other options using an identity policy as described in Create an identity policy, on page 2507.

Step 5 Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

Step 6 Deploy your identity and access control policies to managed devices as discussed in Deploy configuration changes, on page 202.

What to do next

Monitor VPN user traffic .

Hybrid Entra ID and on-premises AD deployments

The hybrid Entra ID and on-premises AD deployment model enables secure, flexible authentication and authorization for remote access users, provided the correct policy order and rule configuration are maintained.

In a hybrid deployment, remote access VPN users can authenticate with Microsoft Entra ID (formerly Azure AD). Cisco ISE authorizes the session and Firewall Management Center uses an on-premises Microsoft Active Directory realm for user identity and user-based access control. In this deployment model:

- Remote access VPN (RA VPN) is configured on Firewall Threat Defense
- Microsoft Entra ID or Azure is used as the SAML identity provider for VPN authentication
- Cisco ISE is used for session authorization
- Cisco ISE authorization policies may include both Microsoft Entra ID/Azure and on-premises AD rules or conditions
- Firewall Management Center uses the on-premises AD domain as the realm for user identity and user enforcement

Best practice: ISE authorization policy order for hybrid deployments

Beginning with Cisco ISE Release 3.4 Patch 4, when ISE evaluates an Entra ID identity source during authorization, it associates Entra ID attributes with the session. If Firewall Management Center uses an on-premises AD realm, the presence of Entra ID attributes may prevent correct user-to-IP mapping for RA VPN sessions. Identity-based access control rules that depend on user or group identity might not match as expected.

To prevent this issue:

Place ISE authorization rules that use the on-premises AD external identity source before rules that evaluate Entra ID or Azure conditions for the same users. This ensures that the session attributes published to FMC through pxGrid align with the configured on-premises AD realm.

After configuring the authorization policy order, verify that new RA VPN sessions create user mappings in Firewall Management Center and that user-based or group-based access control rules match as expected.



Note This behavior applies to Cisco ISE Release 3.4 Patch 4 and later. If user mappings stopped appearing after an ISE upgrade, review the ISE authorization policy order as described in this section.

Troubleshooting guidelines for remote access VPN identity source

For other related troubleshooting information, see Troubleshoot realms and user downloads, on page 2342 and Troubleshoot user control, on page 2526.

If you experience issues with Remote Access VPN, check the connection between your Cloud-Delivered Firewall Management Center and a managed device. If the connection fails, all Remote Access VPN logins

reported by the device cannot be identified during the downtime, unless the users were previously seen and downloaded to the Cloud-Delivered Firewall Management Center.

The unidentified users are logged as Unknown users on the Cloud-Delivered Firewall Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.

The host name of the managed device must be less than 15 characters for Kerberos authentication to succeed.

Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 95

Not Observing Correct Settings for VPN Statistics

This task ensures that managed devices receive the correct health policy settings after modifying VPN statistics configurations.

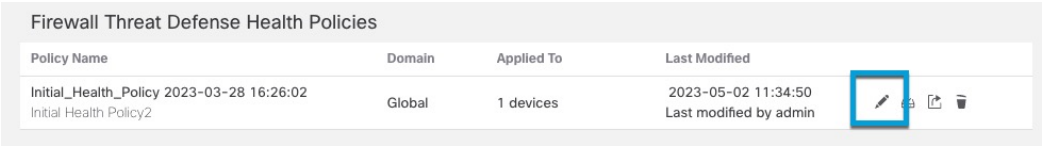
This task discusses steps you must take after either enabling or disabling the **VPN Statistics** setting in a health policy. Failure to perform this task means managed devices have a health policy with incorrect settings.

Procedure

Step 1 Log in to the Cloud-Delivered Firewall Management Center if you haven't already done so.

Step 2 Click **System** (⚙) > **Health** > **Policy**.

Step 3 Under Firewall Threat Defense Health Policies, click **Edit** (✎) next to a policy to edit.



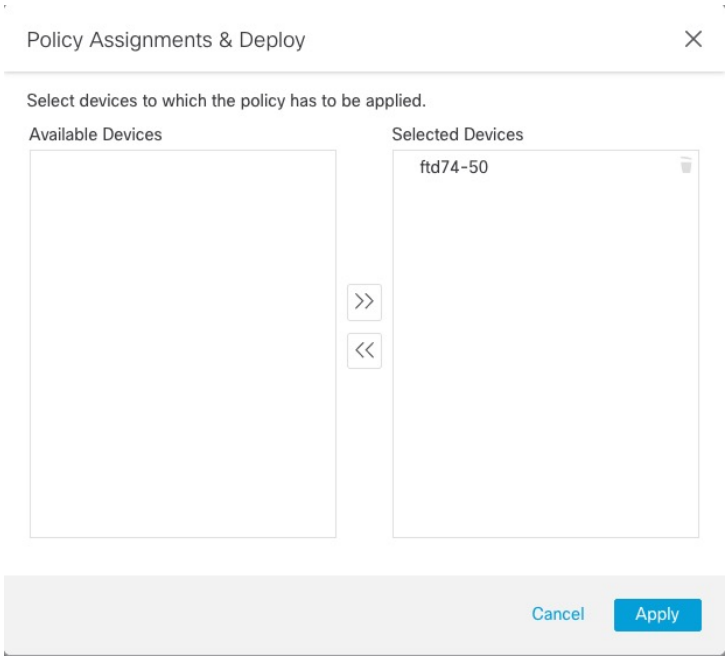
Step 4 On the **Health Modules** tab page, scroll down to locate **VPN Statistics**.

Step 5 Verify the VPN statistics setting is correct or change it if necessary.

Step 6 If you changed the setting, click **Save**, then click **Cancel** to return to the health policy.

Step 7 Under Firewall Threat Defense Health Policies, click **Deploy health policy** (⬇) to apply the policy.

Step 8 In the **Policy Assignments & Deploy** dialog box, move the devices to which to deploy the health policy to the **Selected Devices** field.



Step 9 Click **Apply**.
A message is displayed when the health policy is deployed.

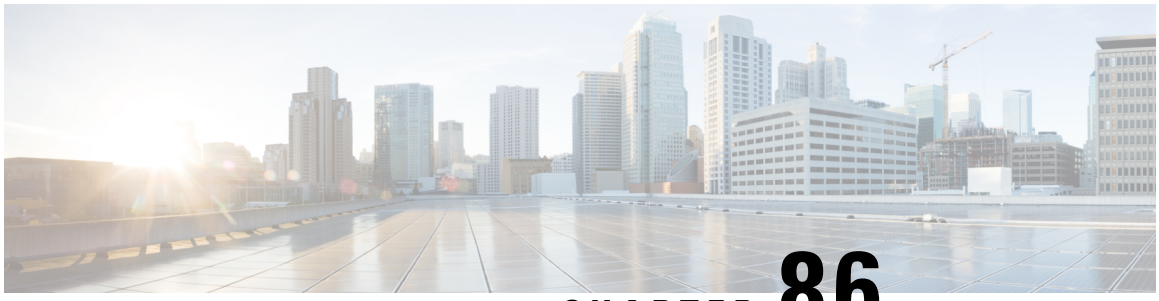
Step 10 After the health policy has finished deploying, click **Policies** > **Access Control heading** > **Access Control** to edit an access control policy.

Step 11 Click **Edit** (✎) next to a policy to edit.

Step 12 Make a minor change to the policy, such as changing its name.

Step 13 Save the access control policy.

Step 14 Deploy configuration changes..



CHAPTER 86

Identity source: TS agent

To use the TS Agent as an identity source for user awareness and user control, install and configure the TS Agent software as discussed in the Terminal Services (TS) Agent guides.

What to do next:

- Specify users to control and other options using an identity policy as described in Create an identity policy, on page 2507.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Deploy your identity and access control policies to managed devices as discussed in Deploy configuration changes, on page 202.
- Monitor user activity.

- Terminal services agent identity source, on page 2501
- Best practices for TS agent configuration, on page 2502
- Identity source: TS agent, on page 2502
- Troubleshoot the TS agent identity source, on page 2502
- History for TS agent, on page 2503

Terminal services agent identity source

The TS Agent is a passive authentication method and one of the authoritative identity sources supported by the system. A Windows Terminal Server performs the authentication, and the TS Agent reports it to a standalone or high availability .

TS agent functionality

When installed on Windows Terminal Servers, the TS Agent assigns a unique port range to individual users as they log in or log out of a monitored network. The Cloud-Delivered Firewall Management Center uses the unique port to identify individual users in the system. You can use one TS Agent to monitor user activity on one Windows Terminal Server and send encrypted data to a Cloud-Delivered Firewall Management Center.

The TS Agent does not report failed login attempts. The data gained from the TS Agent can be used for user awareness and user control.

Best practices for TS agent configuration

Configure the TS Agent on a Windows Terminal Server to monitor user activity and identity realms. The TS Agent enables user awareness in the Cloud-Delivered Firewall Management Center.

- Install and configure the TS Agent on a Windows Terminal Server.
- Configure one or more identity realms targeting the users your server is monitoring.

TS Agent data is visible in the Users, User Activity, and Connection Event tables and can be used for user awareness and user control. For detailed information about the multi-step TS Agent installation and configuration and a complete discussion of the server and system requirements, see the Terminal Services (TS) Agent guides.



Note If the TS Agent monitors the same users as another passive authentication identity source (ISE/ISE-PIC), the Cloud-Delivered Firewall Management Center prioritizes the TS Agent data. If the TS Agent and another passive identity source report activity by the same IP address, only the TS Agent data is logged to the Cloud-Delivered Firewall Management Center.

Identity source: TS agent

To use the TS Agent as an identity source for user awareness and user control, install and configure the TS Agent software as discussed in the Terminal Services (TS) Agent guides.

What to do next:

- Specify users to control and other options using an identity policy as described in Create an identity policy, on page 2507.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- Deploy your identity and access control policies to managed devices as discussed in Deploy configuration changes, on page 202.
- Monitor user activity.

Troubleshoot the TS agent identity source

This reference provides troubleshooting guidance for TS Agent integration issues with the Cloud-Delivered Firewall Management Center. It addresses synchronization requirements, user data prioritization, and activity logging for overlapping passive identity sources.

For other related troubleshooting information, see Troubleshoot realms and user downloads, on page 2342 and Troubleshoot user control, on page 2526.

If you experience issues with the TS Agent integration, check:

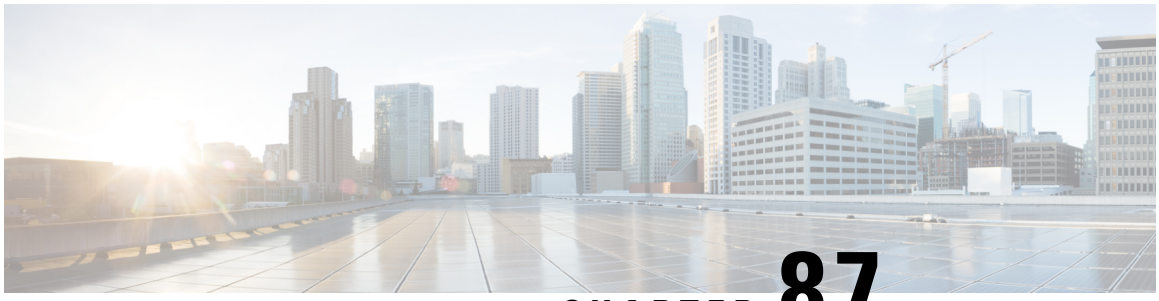
- You must synchronize the time on your TS Agent server with the time on the Cloud-Delivered Firewall Management Center.
- If the TS Agent monitors the same users as another passive authentication identity source (ISE/ISE-PIC), the Cloud-Delivered Firewall Management Center prioritizes the TS Agent data. If the TS Agent and a passive identity source report activity by the same IP address, only the TS Agent data is logged to the Cloud-Delivered Firewall Management Center.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, refer to RFC 959.

For more troubleshooting information, see the Terminal Services (TS) Agent guides.

History for TS agent

This reference provides the historical development timeline and version information for the TS Agent feature, including compatibility details and feature introductions across different Firepower Management Center and Threat Defense versions.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
TS Agent communicates with the Security Cloud Control	Any	7.2.0	By applying a token from Security Cloud Control, the TS Agent can get user login sessions in the same way as for the Security Cloud Control.
TS Agent for user control.	Any	6.2.0	Feature introduced. Firepower now provides the ability to better identify individual users in shared environments, such as Citrix's Virtual Desktop Infrastructure (VDI), to accurately enforce user-based policy rules on the firewall. Users are identified by ports used. The TS Agent software is updated independently of the Firepower Management Center. For more information, see: • *Cisco Terminal Services (TS) Agent Guide* available on cisco.com • Cisco Firepower Compatibility Guide



CHAPTER 87

Identity Policies

The following topics discuss how to create and manage identity rules and identity policies:

- Identity policies, on page 2505
- License requirements for identity policies, on page 2506
- Requirements and prerequisites for identity policies, on page 2506
- Create an identity policy, on page 2507
- Identity rule conditions, on page 2509
- Create an identity rule, on page 2515
- Sample identity policies and rules, on page 2518
- Manage an identity policy, on page 2524
- Manage an identity rule, on page 2525
- Troubleshoot user control, on page 2526

Identity policies

Identity policies contain identity rules. Identity rules associate sets of traffic with a realm and an authentication method: passive authentication, active authentication, or no authentication.

Configuration requirements

With the exception noted below, you must configure realms and authentication methods you plan to use before you can invoke them in your identity rules:

- You configure realms outside of your identity policy, at **Integration** > **Other Integrations** > **Realms** > **Realms**. For more information, see Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
- You configure ISE/ISE-PIC, a passive authentication identity source, at **Integration** > **Other Integrations** > **Identity Sources**. For more information, see Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source, on page 2390.
- You configure the TS Agent, a passive authentication identity source, outside the system. For more information, see the *Cisco Terminal Services (TS) Agent Guide*.
- You configure captive portal, an active authentication identity source, within the identity policy. For more information, see Configure the captive portal for user control, on page 2355.

- You configure Remote Access VPN, an active authentication identity source, in Remote Access VPN policies. For more information, see Configure AAA settings for a remote access VPN policy.

After you add multiple identity rules to a single identity policy, order the rules. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles the traffic.

You can also filter traffic by network object, which limits the network each device monitors in the event your devices are at or near their memory limits.

After you configure one or more identity policies, you must associate one identity policy with your access control policy. When traffic on your network matches the conditions in your identity rule, the system associates the traffic with the specified realm and authenticates the users in the traffic using the specified identity source.

If you do not configure an identity policy, the system does not perform user authentication.

Exception to creating an identity policy

An identity policy is not required if all these conditions are true:

- You use the ISE/ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies.

Related Topics

Set up an identity policy, on page 2263

License requirements for identity policies

This reference provides license requirements for identity policies.

Threat Defense License

Any

Classic License

Control

Requirements and prerequisites for identity policies

This reference provides the requirements and prerequisites necessary for implementing identity policies.

Model support

Any.

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Create an identity policy

Create an identity policy to control user and group access in a realm for access control policies.

This task discusses how to create an identity policy.

Before you begin

An identity policy is required to use users and groups in a realm in access control policies. Create and enable one or more realms as described in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

(Optional.) If a particular managed device monitors a large number of user groups, the system might drop user mappings based on groups due to managed device memory limitations. As a result, rules with realm or user conditions might not perform as expected. Provided the devices run version 6.7 or later, you can configure the identity rule to monitor traffic by one network or network group object only. To create a network object, see Create a network object, on page 2985.

An identity policy is not required if all these conditions are true:

- You use the ISE/ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** and click **New Policy**.

Step 3 Enter a **Name** and, optionally, a **Description**.

Step 4 Click **Save**.

Step 5 To add a rule to the policy, click **Add Rule** as described in Create an identity rule, on page 2515.

Step 6 To create a rule category, click **Add Category**.

Step 7 To configure captive portal active authentication, click **Active Authentication** and see Create an identity policy and active authentication rule, on page 2358.

Step 8 (Optional.) To filter traffic by network object, click the **Identity Source** tab. From the list, click the network object to use to filter traffic for this identity policy. Click **Add** (+) to create a new network object.

Step 9 Click **Save** to save the identity policy.

What to do next

- Add rules to your identity policy that specify which users to match and other options; see Create an identity rule, on page 2515.
- Associate the identity policy with an access control policy to allow or block selected users from accessing specified resources; see Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- (Not required for Microsoft Azure AD realms.) Deploy configuration changes to managed devices; see Deploy configuration changes, on page 202.

If you encounter issues, see Troubleshoot user control, on page 2526.

Related Topics

Create an identity policy and active authentication rule, on page 2358
Captive portal fields, on page 2367
Troubleshoot user control, on page 2526
Create an identity mapping filter, on page 2508

Create an identity mapping filter

An identity mapping filter can be used to limit the networks to which an identity rule applies. For example, if your Firewall Management Center manages FTDs that have a limited amount of memory, you can limit the networks they monitor.

You must create separate identity mapping filters for IPv4 and IPv6 addresses.

You can also optionally exclude subnets from the following:

- Receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE.

You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.

Before you begin

Perform the following tasks:

1. Create a realm, which is required for an identity policy. See Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.
2. Create an identity policy. See Create an identity policy, on page 2507.
3. Create a network object or network group object as discussed in Create a network object, on page 2985. The network object or group you create should define the network you want managed devices to monitor in identity policies.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity**.

Step 3 Click **Edit** ( ).

Step 4 Click the **Identity Source** tab.

Step 5 From the **Identity Mapping Filter** list, choose the name of a network object to use as a filter .

To create a new network object, see Create a network object, on page 2985.

Note
To restrict traffic to IPv6 addresses, you must add at least one address, network, or group to the filter.

Step 6 Click **Save**.

Step 7 (Not required for Microsoft Azure AD realms.) Deploy configuration changes to managed devices; see Deploy configuration changes, on page 202.

What to do next

(Not required for Microsoft Azure AD realms.) Associate the identity policy with an access control policy as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

To check or change ISE identity mapping filters (also referred to as *subnet filters*), use these commands:

```
show identity-subnet-filter
configure identity-subnet-filter { add | remove } subnet
```

Identity rule conditions

Identity rule conditions enable you to fine-tune your identity policy to target specific users and networks you want to control. See one of the following topics for more information.

Related Topics

Security zone rule conditions, on page 2202
Network rule conditions, on page 1701
VLAN tag rule conditions, on page 2070
Best practice: use application-based rules over port-based rules, on page 1703
Realm & settings rule conditions, on page 2513

Security zone rule conditions

A security zone rule condition is a network traffic control mechanism that

- segments your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices,
- controls or decrypts traffic by its source and destination security zones, and
- requires all interfaces in a zone to be of the same type (all inline, passive, switched, or routed).

Zone condition traffic matching

If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

All zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.



Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Network rule conditions

A network condition is a traffic matching criterion that

- controls or decrypts traffic based on source and destination IP addresses using inner packet headers,
- enables administrators to specify exact IP addresses or blocks for granular policy enforcement, and
- can be built using predefined network objects or by manually entering IP addresses or address blocks.

Network rule condition restrictions

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.



Note You *cannot* use FDQN network objects in identity rules.

Redirect to host name network rule conditions

(Snort 3.0 only)—You can use a network object that contains the fully-qualified host name (FQDN) of the interface that captive portal can use for active authentication requests.

The FQDN must resolve to the IP address of one of the interfaces on a managed device. By using an FQDN, you can assign a certificate for active authentication that the client will recognize, thus avoiding the untrusted certificate warning users get when being redirected to a managed device's IP address.

Certificate and redirection behavior

The certificate can specify one FQDN, a wildcard FQDN, or multiple FQDNs in the Subject Alternate Names (SAN) in the certificate.

If an identity rule requires active authentication for a user, but you do not specify a redirect FQDN, the user is redirected to the captive portal port on the managed device interface to which they are connected.

If you do not supply a Redirect to Host Name FQDN, the HTTP Basic, HTTP Response Page, and NTLM authentication methods redirect the user to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.directory-server-domain-name*. If you use HTTP Negotiate without a Redirect to Host Name FQDN, update your DNS server to map this name to the IP addresses of all inside interfaces that require active authentication. If you do not update the DNS server, the redirection does not complete and users cannot authenticate.

We recommend that you always provide a Redirect to Host Name FQDN to ensure consistent behavior regardless of authentication method.

VLAN tag rule conditions

VLAN tags rule conditions are filtering mechanisms that

- control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic
- use the innermost VLAN tag to filter VLAN traffic, with the exception of prefilter policy which uses the outermost VLAN tag, and
- apply only to inline sets and do not match traffic on firewall interfaces.

VLAN tag rule conditions configuration details



Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note this Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

Best practice: use application-based rules over port-based rules

Use application filtering criteria instead of port filtering to target traffic. Applications can be configured to use unique ports to bypass access control blocks.

Port conditions allow you to control traffic by its source and destination ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

The traditional method targets applications by using port specifications, but applications can manipulate unique ports to bypass these controls

Application filtering is recommended for applications to avoid blocking legitimate applications, like Firewall Threat Defense, that dynamically open separate channels for control and data flow. Port-based access control rules could block such applications, impacting desirable operations.

Using source and destination port constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP).

- For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).
- If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Port, protocol, and ICMP code rule conditions (concept)

A port, protocol, and ICMP code rule condition is a traffic matching criterion that

- matches traffic based on source and destination ports
- supports TCP, UDP, ICMP, and other protocols that do not use ports, and
- controls traffic flow in access control and other policy rules.

Port condition types

Port conditions match traffic based on the source and destination ports. Depending on the rule type, "port" can represent any of these types:

- TCP and UDP—Control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP (6) /22.
- ICMP—Control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—Control traffic using other protocols that do not use ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best practices for port-based rules

Specifying ports is the traditional method used to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Use application filtering criteria when possible; however, note that application filtering is unavailable in prefilter rules.

Use application filtering for FTP and similar applications that open separate channels for control and data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly and could result in blocking desirable connections.

Using source and destination port constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching non-TCP traffic with port conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although configuring port conditions can match non-TCP traffic, restrictions apply:

- Access control rules—For Classic devices, match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, add only network-based conditions: zone, IP address, port, and VLAN tag. The system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support only TCP port conditions.
- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies, sent in response to ICMP echo requests, are ignored. To match any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Realm & settings rule conditions

The **Realm & Settings** tab page enables you to choose a realm or realm sequence to which to apply the identity rule. If you are using captive portal, you have additional options.

Authentication realm

From the **Realm** list, click a realm or realm sequence.

The realm or realm sequence containing the users you want to perform the specified **Action** on. You must fully configure a realm or realm sequence before selecting it as the realm in an identity rule.



Note If remote access VPN is enabled and your deployment is using a RADIUS server group for VPN authentication, make sure you specify the realm associated with this RADIUS server group.

Active authentication only: other options

If you either choose **Active Authentication** as the authentication type or if you check the box, **Use active authentication if passive or VPN identity cannot be established**, you have these options.

Use active authentication if passive or VPN identity cannot be established

(Passive authentication rule only.) Selecting this option authenticates users using captive portal active authentication if a passive or a VPN authentication fails to identify them. You must configure an Active Authentication rule in your identity policy in order to select this option. (That is, users must authenticate using the captive portal.)

If you disable this option, users that do not have a VPN identity or that passive authentication cannot identify are identified as Unknown.

Also see the discussion of the **Authentication Realm** list later in this topic,

Identify as Special Identities/Guest if authentication cannot identify user

Selecting this option allows users who fail captive portal active authentication the specified number of times to access your network as a guest. These users appear in the Cloud-Delivered Firewall Management Center identified by their username (if their username exists on the AD or LDAP server) or by **Guest** (if their user name is unknown). Their realm is the realm specified in the identity rule. (By default, the number of failed logins is 3.)

This field is displayed only if you configure **Active Authentication** (that is, captive portal authentication) as the rule **Action**.

Authentication Protocol

The method to use to perform captive portal active authentication. .

The selections vary depending on the type of realm, LDAP or AD:

- Choose **HTTP Basic** if you want to authenticate users using an unencrypted HTTP Basic Authentication (BA) connection. Users log in to the network using their browser's default authentication pop-up window.

  Most web browsers cache the credentials from **HTTP Basic** logins and use the credentials to seamlessly begin a new session after an old session times out.

- Choose **NTLM** to authenticate users using a NT LAN Manager (NTLM) connection. This selection is available only when you select an AD realm. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication pop-up window.

- Choose **Kerberos** to authenticate users using a Kerberos connection. This selection is available only when you select an AD realm for a server with secure LDAP (LDAPS) enabled. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication pop-up window.

Note The **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.

Note If you are creating an identity rule to perform Kerberos captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the host name you provided when configuring DNS.

For Firewall Threat Defense devices, the FQDN must resolve to the IP address of the routed interface used for captive portal.

- Choose **HTTP Negotiate** to allow the captive portal server to choose between HTTP Basic, Kerberos, or NTLM for the authentication connection. This type is available only when you select an AD realm.

Note The **Realm** you choose must be configured with an **AD Join Username** and **AD Join Password** for **HTTP Negotiate** to choose Kerberos captive portal active authentication.

Note If you are creating an identity rule to perform **HTTP Negotiate** captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN of the device you are using for captive portal must match the hostname you provided when configuring DNS.

- Choose **HTTP Response Page** to enable users to choose a realm to log in to.

  You can optionally customize the response page; for example, to conform to company style standards.

Create an identity rule

Identity rules enable you to control user authentication and traffic handling by specifying which identity realm should authenticate users for specific network criteria.

For details about configuration options for identity rules, see Identity rule fields, on page 2517.

Before you begin

You must create and enable a realm or realm sequence.

- Create a Microsoft Active Directory realm and realm directory as discussed in Create an LDAP realm or an Active Directory realm and realm directory, on page 2325.

- (Microsoft AD realm only.) Download users and groups and enable the realm as discussed in Synchronize users and groups, on page 2330.
- Create a Microsoft Azure AD (SAML) realm as discussed in Create a Microsoft Azure AD (SAML) realm, on page 2289.
- (Optional.) Create a realm sequence as discussed in Create a realm sequence, on page 2332.
- Rules are evaluated top-down. For a connection that matches the specified network criteria of a given rule, the user is evaluated against the identity realm specified in the rule. If the user is not part of that realm, they will be marked as unknown, and no further rules in the identity policy will be evaluated. Thus, if you have more than one realm that needs to be evaluated, be sure to use realm sequences instead of a single realm.



Caution Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include a decryption policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity**.

Step 3 Click **Edit** ( ) next to the identity policy to which to add the identity rule.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **Add Rule**.

a) Enter a **Name**.
b) If the Specified rule is applicable, check the check box of **Enabled**.
c)
d) To add the rule to an existing category, indicate where you want to **Insert** the rule. To add a new category, click **Add Category**.
e) Choose a rule **Action** from the list.
f) If you're configuring captive portal, see Configure the captive portal for user control, on page 2355.
g) (Optional) To add conditions to the identity rule, see Identity rule conditions, on page 2509.
h) Click **Add**.

Step 5 In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste. Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number.

The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes.

Identity rule fields

Configure identity rules to specify authentication actions for users, including Passive Authentication, Active Authentication, or No Authentication. Learn how VPN identity preference affects authentication and rule processing in remote access scenarios.

Use these fields to configure identity rules.

Enabled

Enabling this option enables the identity rule in the identity policy. Unselecting this option disables the identity rule.

Action

Specify the type of authentication you want to perform on the users in the specified realm: **Passive Authentication** (default), **Active Authentication**, or **No Authentication**. You must fully configure the authentication method, or *identity source*, before selecting it as the action in an identity rule.

Additionally, if VPN is enabled (configured on at least one managed device), remote access VPN sessions are actively authenticated by VPN. Other sessions use the rule action. This means that, if VPN is enabled, VPN identity determination is performed first for all sessions regardless of the selected action. If a VPN identity is found on the specified realm, this is the identity source used. No additional captive portal active authentication is done, even if selected.

If the VPN identity source is not found, the process continues according to the specified action. You cannot restrict the identity policy to VPN authentication only because if the VPN identity is not found, the rule is applied according to the selected action.



Caution Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include an SSL policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

For information about which passive and active authentication methods are supported in your version of the system, see About user identity sources, on page 2256.

Sample identity policies and rules

These sections provide examples of configuring an identity policy with either a passive authentication rule or an active authentication rule. In addition, because you can authenticate users with active authentication using either realms or realm sequences, separate examples are provided.

Authentication types

Active authentication means the user authenticates using the captive portal; the user enters network credentials before being allowed to access allowed resources. (RA-VPN is another type of active authentication but it cannot be used together with captive portal authentication. For more information, see Remote access VPN identity sources, on page 2495.)

Passive authentication covers all other types. Passive authentication includes using Microsoft Active Directory realms, Microsoft Azure Active Directory realms, Cisco Identity Services Engine, and others.

These examples use these assumptions:

- Microsoft Active Directory (AD) realm named forest.example.com with two child domains configured in a trust relationship:
  - US-West
  - US-East
- A realm sequence named US that includes both realms
- Passive authentication rule that authenticates users with a realm sequence
- Two active authentication rules:
  - One rule that authenticates users with a realm and uses the NTLM authentication protocol
  - One rule that authenticates users with a realm sequence and uses the HTTP Response Page authentication protocol
- Each sample identity rule is associated with a different identity policy

When you configure a Passive Authentication identity rule, you can choose to authenticate users either with LDAP, a Microsoft Active Directory realm, or a Microsoft AD realm sequence. You can use a realm to authenticate with any authentication type; realm sequence restricts the authentication types you can use. For an example, see Create an identity policy with a passive authentication rule, on page 2519.

When you configure an Active Authentication identity rule, you can choose to authenticate users either with LDAP, a Microsoft Active Directory realm, or a Microsoft AD realm sequence. You can use a realm to authenticate with any authentication type; realm sequence restricts the authentication types you can use.

You can also authenticate users with a Microsoft Active Directory realm sequence except for these authentication types:

- NTLM
- Kerberos
- HTTP Negotiate

For an example, see Create a sample identity policy with an active authentication rule, on page 2520.

Create an identity policy with a passive authentication rule

Create an identity policy that authenticates users using passive authentication with realm sequences for user identification.

This task discusses how to create an identity policy with a passive authentication rule that authenticates users using the US realm sequence. If a user is not found in the first realm in the sequence, the system searches the other realms in the sequence in the order listed in the realm sequence. If a user is still not found in the realm or realm sequence, the user is identified as Unknown.

You can optionally authenticate a user with the captive portal (that is, active authentication) if the user is not found in any realm in the sequence. For more information, see Best practices for captive portal guidelines and limitations, on page 2353.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity**.

Step 3 Click **New Policy**.

a) Enter a **Name** for the policy and an optional **Description**.

b) Click **Save**.

Step 4 Click **Add Rule**.

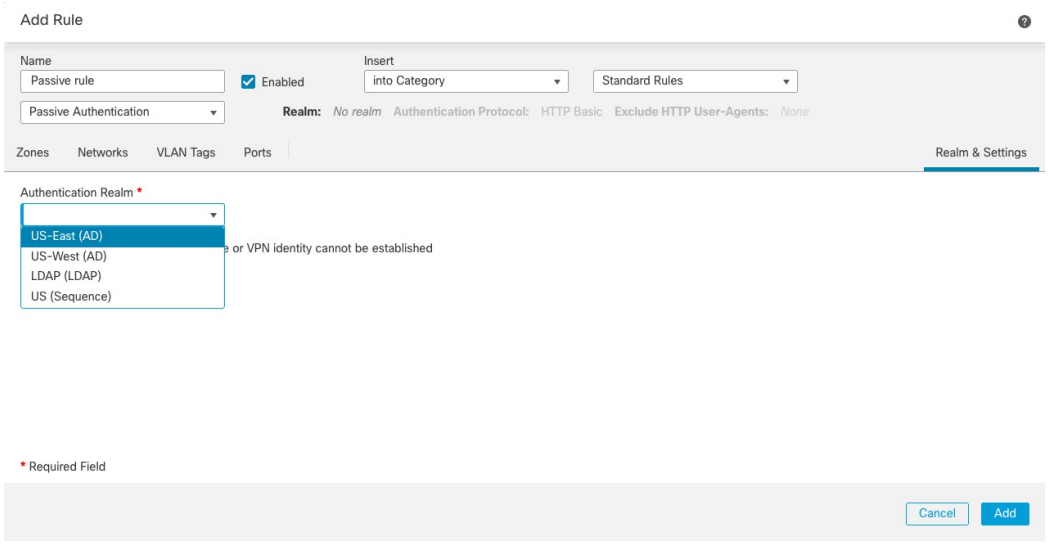
a) Enter a **Name** for the rule.

b) From the list, click **Passive Authentication**.

Step 5 Click the **Realms & Settings** tab page.

a) From the list, click the name of a realm or realm sequence.

The following figure shows an example.



- If you choose a realm (such as **US-East** in the example), the system searches that realm for users to match the rule. If a user is not found, the user is identified as Unknown.
- If you choose a realm sequence (**US (Sequence)** in the example), users are searched in every realm in the sequence in the order specified in the sequence. If the user is not found, the user is identified as Unknown.
- You can also choose an LDAP realm.
- For additional ways to authenticate the user, check **Use active authentication if passive or VPN identity cannot be established**. For more information, see Best practices for captive portal guidelines and limitations, on page 2353.

Step 6 (Optional.) To filter traffic by network object, click the **Identity Source** tab. From the list, click the network object to use to filter traffic for this identity policy. Click **Add** (+) to create a new network object.

Step 7 Set identity rule conditions as discussed in Identity rule conditions, on page 2509.

Step 8 Associate the identity rule with an access control rule as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

Step 9 Deploy configuration changes to managed devices; see Deploy configuration changes, on page 202.

Create a sample identity policy with an active authentication rule

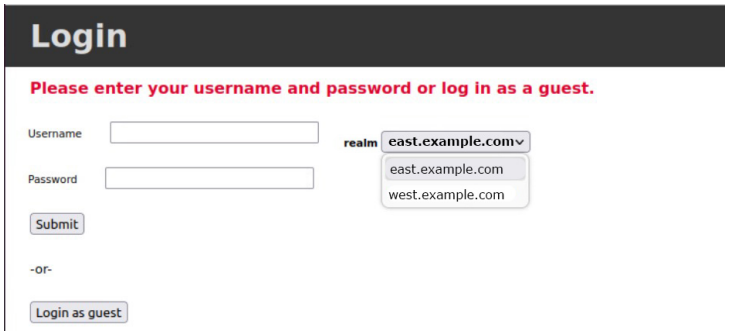
This task shows you how to configure an identity policy with an Active Authentication rule where authentication is performed using either a realm or a realm sequence. This enables secure network access management through customizable captive portal settings.

The associated tasks show an examples of configuring an identity policy with an **Active Authentication** rule where authentication is performed using either a realm or a realm sequence.

The differences follow:

- *Realm* enables you to use any supported authentication type (currently, **HTTP Basic**, **NTLM**, **Kerberos**, **HTTP Negotiate**, or **HTTP Response Page**).
- *Realm sequence* limits you to only the **HTTP Basic** or **HTTP Response Page** authentication types.

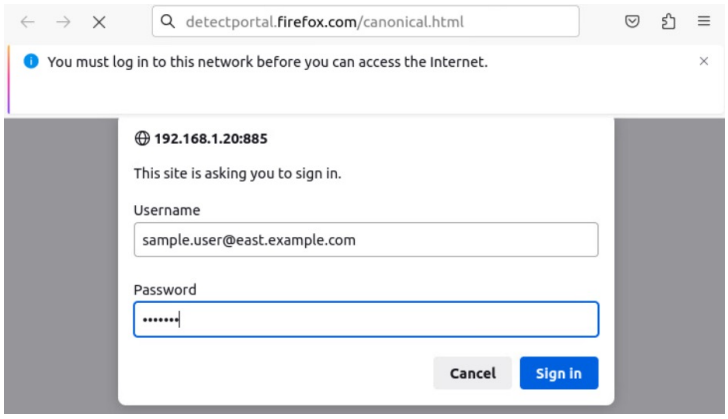
  Users who authenticate with a realm sequence and the HTTP Response Page authentication type see the following by default:



  The user can authenticate in any of the following ways:

- If a list of realms in the realm sequence is displayed (as shown), the user must enter their user name and password in the provided fields and click the name of their realm from the list.
- If realms are not displayed in a list, the user can enter their credentials in `username@domain` format.

Users who authenticate with a realm and the HTTP Basic authentication page see the following:



The user must enter their user name in `username@domain` format.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity**.

Step 3 Click **New Policy**.

a) Enter a **Name** for the policy and an optional **Description**.

b) Click **Save**.

Step 4 Click the **Active Authentication** tab.

Enter the following information:

- **Server Certificate**: From the list, click an internal certificate object to use for secure connection to the Firewall Threat Defense device, or click **Add** (+) to add one.
- **Redirect to Host Name**: (Optional.) From the list, click a network object to which to redirect captive portal requests. If you omit this value, requests redirect to the IP address of the managed device. You can click **Add** (+) to create a new network object. For more information, see Redirect to host name network rule conditions, on page 2510.

  Managed devices must have Snort 3 enabled to use this option.
- **Port**: Enter a port for the captive portal to use. This port must be unique for the captive portal and must match an access control rule you set up as discussed in Create a TCP port access control rule, on page 2360. (Default is 885.)
- **Maximum login attempts**: Enter the maximum number of login attempts before login fails. (Default is 3.)

- **Active Authentication Response Page**: Choose a system-provided or custom login page for captive portal users. For more information about your options, see Captive portal fields, on page 2367.

Step 5 Click **Save** to save your changes to the identity policy.

Step 6 Click the **Rules** tab.

Step 7 Click **Add Rule**.

a) Enter a **Name** for the rule.

b) From the list, click **Active Authentication**.

Step 8 Click the **Realms & Settings** tab page. Then configure an active authentication method for the users.

What to do next

Continue with one of the following sections:

- Active authentication using a realm, on page 2522
- Active authentication using a realm sequence, on page 2523

Active authentication using a realm

This task authenticates captive portal users with a realm and any available authentication protocol (currently, **HTTP Basic**, **NTLM**, **Kerberos**, **HTTP Negotiate**, or **HTTP Response Page**).

Active authentication using a realm provides secure user verification for captive portal access by integrating with your organization's authentication infrastructure.

Before you begin

Complete the task discussed in Create a sample identity policy with an active authentication rule, on page 2520.

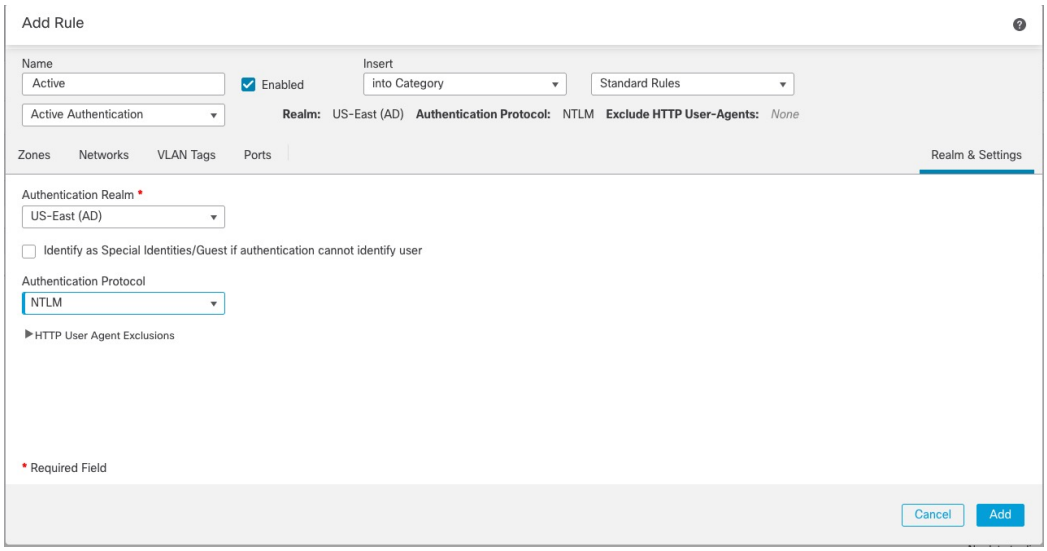
Procedure

Step 1 Continue from Create a sample identity policy with an active authentication rule, on page 2520.

Step 2 On the **Realms & Settings** tab page, click **US-East**.

Step 3 From the **Authentication Protocol** list, click **NTLM**.

The following figure shows an example.



If you choose a realm (such as in the example), the system searches that realm for users to match the rule. If a user is not found, the user is identified as Unknown.

Step 4 Click **Add**.

Step 5 (Optional.) To filter traffic by network object, click the **Identity Source** tab. From the list, click the network object to use to filter traffic for this identity policy. Click **Add** (+) to create a new network object.

Step 6 Set identity rule conditions as discussed in Identity rule conditions, on page 2509.

Step 7 Associate the identity rule with an access control rule as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

Step 8 Deploy configuration changes to managed devices; see Deploy configuration changes, on page 202.

Active authentication using a realm sequence

This task enables authentication of captive portal users using a realm sequence, allowing you to control access through HTTP Basic or HTTP Response Page protocols. It helps you select realms, set identity rule conditions, and associate them with access control rules to filter network traffic and manage user identification.

This task discusses how to authenticate captive portal users with a realm sequence, which restricts you to the **HTTP Basic** or **HTTP Response Page** authentication protocols.

Before you begin

Complete the task discussed in Create a sample identity policy with an active authentication rule, on page 2520.

Procedure

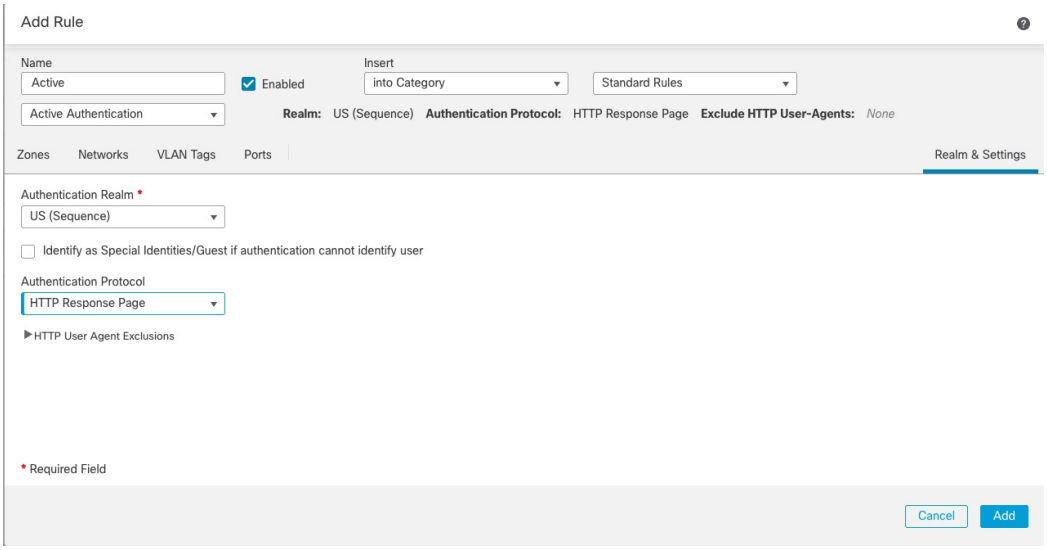
Step 1 Continue from Create a sample identity policy with an active authentication rule, on page 2520.

Step 2 On the **Realms & Settings** tab page, click the name of a realm from the list.

Step 3 From the list, click **US-East**.

Step 4 From the **Protocol** list, click **HTTP Response Page**.

The following figure shows an example.



If you choose a realm sequence (such as in the example), the system searches realms in the sequence in the order specified in the realm sequence. The first realm in the sequence can be referred to as the *default* realm; it's the realm that is used if the user does not change it. If a user is not found, the user is identified as Unknown.

(Only if you upgraded to version 7.4.1 from an earlier version.) Edit the HTTP response page so it shows the list of realms in the sequence as discussed in Update a custom authentication form, on page 2360.

Step 5 Click **Add**.

Step 6 (Optional.) To filter traffic by network object, click the **Identity Source** tab. From the list, click the network object to use to filter traffic for this identity policy. Click **Add** (+) to create a new network object.

Step 7 Set identity rule conditions as discussed in Identity rule conditions, on page 2509.

Step 8 Associate the identity rule with an access control rule as discussed in Associate prefilter, decryption, and identity policies with an access control, on page 2016.

Step 9 Deploy configuration changes to managed devices; see Deploy configuration changes, on page 202.

Manage an identity policy

This task allows you to manage identity policies by performing various administrative actions including editing, deleting, copying, generating reports, and organizing policies into categories for optimized security policy administration.

Identity policies control access to network resources based on user identity. Managing these policies effectively ensures proper security enforcement and administrative organization within the Secure Firewall Management Center.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity**.

Step 3 To delete a policy, click **Delete** ( ). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 To edit a policy, click **Edit** ( ) next to the policy and make changes as described in Create an identity policy, on page 2507. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 5 To copy a policy, click **Copy** ( ).

Step 6 To generate a report for the policy, click **Report** ( ) as described in Generate current policy reports, on page 220.

Step 7 To compare policies, see Compare policies, on page 219.

Step 8 To create a folder in which to organize policies, click **Add Category**.

What to do next

Deploy configuration changes.

Manage an identity rule

This task allows you to edit, delete, or organize identity rules in access control policies to manage user access and security within the Secure Firewall Management Center.

Use this procedure to modify existing identity rules, remove unnecessary rules, or organize rules using categories to maintain effective security policies.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity**.

Step 3 Click **Edit** ( ) next to the policy you want to edit. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 To edit an identity rule, click **Edit** ( ) and make changes as described in Create an identity policy, on page 2507.

Step 5 To delete an identity rule, click **Delete** ( ).

Step 6 To create a rule category, click **Add Category** and choose the position and the rule.

Step 7 Click **Save**.

What to do next

- Deploy configuration changes.

Troubleshoot user control

If you notice unexpected user rule behavior, consider tuning your rule, identity source, or realm configurations. For other related troubleshooting information, see:

- Troubleshoot Cisco ISE/ISE-PIC or Cisco TrustSec issues, on page 2399
- Troubleshoot the TS agent identity source, on page 2502
- Troubleshoot the captive portal identity source, on page 2369
- Troubleshoot realms and user downloads, on page 2342

Rules targeting realms, users, or user groups are not matching traffic

If you configure a TS Agent or ISE/ISE-PIC device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user records due to your Cloud-Delivered Firewall Management Center user limit. As a result, rules with user conditions may not match traffic as expected.

Rules targeting user groups or users within user groups are not matching traffic as expected

If you configure a rule with a user group condition, your LDAP or Active Directory server must have user groups configured. The system cannot perform user group control if the server organizes the users in basic object hierarchy.

Rules targeting users in secondary groups are not matching traffic as expected

If you configure a rule with a user group condition that includes or excludes users who are members of a secondary group on your Active Directory server, your server may be limiting the number of users it reports.

By default, Active Directory servers limit the number of users they report from secondary groups. You must customize this limit so that all of the users in your secondary groups are reported to the Cloud-Delivered Firewall Management Center and eligible for use in rules with user conditions.

Rules are not matching users when seen for the first time

After the system detects activity from a previously-unseen user, the system retrieves information about them from the server. Until the system successfully retrieves this information, activity seen by this user is *not* handled by matching rules. Instead, the user session is handled by the next rule it matches (or the policy's default action, if applicable).

For example, this might explain when:

- Users who are members of user groups are not matching rules with user group conditions.
- Users who were reported by a TS Agent or ISE device are not matching rules, when the server used for user data retrieval is an Active Directory server.

Note that this might also cause the system to delay the display of user data in event views and analysis tools.

Rules are not matching all ISE/ISE-PIC users

This is expected behavior. You can perform user control on ISE/ISE-PIC users who were authenticated by an Active Directory domain controller. You cannot perform user control on ISE/ISE-PIC users who were authenticated by an LDAP, RADIUS, or RSA domain controller.

Users and groups using too much memory

If processing users and groups is using too much memory, health alerts are displayed. Remember that all user sessions are propagated to all devices managed by the Cloud-Delivered Firewall Management Center. If your Cloud-Delivered Firewall Management Center manages devices with different amounts of memory, the device with the least amount of memory determines the number of user sessions the system can handle without errors.

It's not possible to tune memory allocated to identity processes; even if a device has available memory, it can report out-of-memory issues. If issues persist, you have these options:

- Segregate lower capacity managed devices on subnets and configure ISE/ISE-PIC to not report passive authentication data to those subnets.

  See the chapter on managing network devices in the *Cisco Identity Services Engine Administrator Guide*.

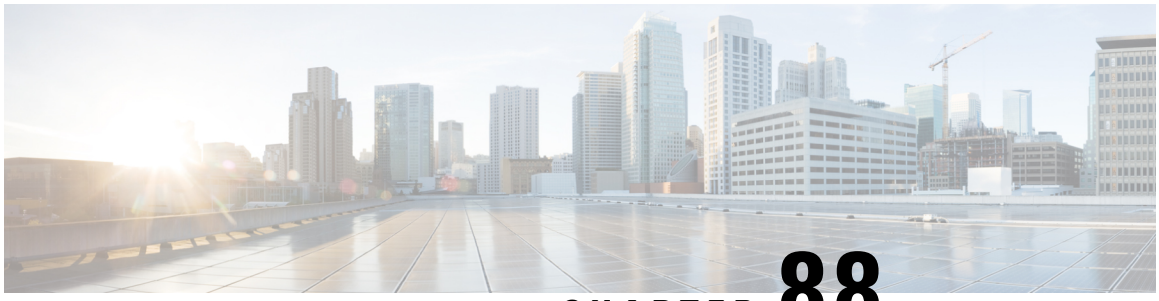
- Unsubscribe from Security Group Tags (SGTs).
- Upgrade your managed device to a model with more memory.



PART XX

Advanced Policies and Settings for Access Control

- Prefilter Policies , on page 2531
- DNS Policies for Security Intelligence, on page 2557
- File Policies for Network Malware Protection, on page 2583
- Dynamic Attributes Connector, on page 2621
- Advanced Settings for Access Control, on page 2687
- Service Policies, on page 2707
- Threat Detection, on page 2725
- Elephant Flow Detection, on page 2735
- Policy Analyzer and Optimizer, on page 2749



CHAPTER 88

Prefilter Policies

The prefilter policy is the first security policy applied to incoming connections. Prefilter rules evaluate traffic on layer 3/4 criteria only, that is, protocol and source/destination IP address and port. They give you a chance to make early decisions on connections so you can avoid further processing and improve device performance.

- Prefilter policies, on page 2531
- Requirements and prerequisites for prefilter policies, on page 2540
- Guidelines and limitations for the prefilter policy, on page 2541
- Guidelines and limitations for encapsulated traffic handling, on page 2541
- Best practices for fastpath prefiltering, on page 2542
- Configure prefilter policies, on page 2543
- Examples of prefilter policy, on page 2552

Prefilter policies

A prefilter policy is a security policy that

- serves as the first security policy applied to incoming connections
- evaluates traffic on layer 3/4 criteria only, including protocol and source/destination IP address and port
- enables early decisions on connections to avoid further processing and improve device performance, and
- can fast path connections and make them eligible for flow offload.

Prefilter policy actions

Use prefilter policies to implement these main actions:

- Improve performance—The sooner you exclude traffic that does not require inspection, the better. Specifically, you can perform these actions.
  - Offload trusted flows—If your device supports flow offload, use the prefilter Fastpath action to identify flows you trust to be eligible for offloading to NIC processing. No further processing, including inspection, is needed for these connections. Refer to Best practices for fastpath prefiltering, on page 2542 and Large flow offloads, on page 2535.
  - Block unwanted connections—For example, if you want to block entire subnets, protocols, or TCP/UDP ports, configure this in the prefilter policy.

- Rezone plain text encapsulated tunnels so you can handle them as a unit in the access control policy. Some protocols, such as GRE, create plain-text tunnels that include many connections. By rezoning a tunnel, you assign it a tag, which you can select in access control rules as a source or destination zone. This allows you to apply customized intrusion or other rules to the tunneled traffic and provides granular insight and control.

Note VPN or other encrypted tunnels are not plain-text tunnels and cannot be rezoned.

You must assign each access control policy a prefilter policy. However, the system comes with a pre-defined prefilter policy that might work for you.

Default prefilter policy

A default prefilter policy only analyzes plain-text encapsulated tunnels. It is automatically assigned when you create a new access control policy. It cannot be deleted from the system.

Default prefilter policy characteristics

Every access control policy must have an assigned prefilter policy. For convenience, the system includes a default prefilter policy that is automatically assigned when you create a new access control policy.

You can instead change the default policy to block all plain-text tunnels. See Configuring the default action, on page 2545.

If you want to do any other customization, including adding prefilter or tunnel rules, you must create your own prefilter policy and assign it to the appropriate access control policies, as explained in Assign access control policy to devices, on page 2014.

Prefilter policy rule order

Prefilter policy rule order is a traffic processing mechanism that determines which tunnel rules and prefilter rules handle network traffic first and applies the default action for unmatched tunnels.

Rule processing sequence

In a prefilter policy, tunnel rules, prefilter rules, and a default action handle network traffic:

- Tunnel and prefilter rules—First, rules in a prefilter policy handle traffic in the order you specify. Tunnel rules match specific tunnels only. Prefilter rules have a wider range of constraints.

  Rules are match in top-down order, first match wins. Thus, rule order is critical. For example, you might want to block an entire subnet except for two specific IP addresses on the subnet. In that case:

  - The first rule should specify the two allowable IP addresses as the source network, and the rule action would be Analyze. Connections from these two IP addresses would match this rule, and the next rule would not be processed.
  - The next rule should specify the entire subnet as the source network, with Block for the rule action.

  For more information, see Tunnel vs prefilter rules, on page 2533.

- Default action (tunnels only)—If a tunnel does not match any rules, the default action handles it. The default action can block these tunnels or continue access control on their individual encapsulated connections (Analyze action).

  There is no default action for nonencapsulated traffic. If a nonencapsulated connection does not match any prefilter rules, the system continues with access control. These connections might subsequently be blocked by other policies, or ultimately allowed.

If you need to change rule order, you can:

- Drag and drop the rule to the right location.
- Right click the rule and **Cut**, then find the right location, right click, and select **Paste Above** or **Paste Below**.
- Edit the rule and use the **Insert** control to change its location.

Tunnel vs prefilter rules

Choose to configure a tunnel rule or a prefilter rule based on the type of traffic you want to match and the actions or analysis you intend to perform.

Characteristic	Tunnel Rules	Prefilter Rules
Primary function	Quickly fastpath, block, or rezone plaintext, passthrough tunnels.	Quickly fastpath or block any other connection that benefits from early handling.
Encapsulation and port/protocol criteria	Encapsulation conditions match only plaintext tunnels over selected protocols: GRE, IP-in-IP, IPv6-in-IP, Teredo.	Port conditions can use a wider range of port and protocol constraints than tunnel rules; see Port, protocol, and ICMP code rule conditions (concept), on page 1703.
Network criteria	Tunnel endpoint conditions constrain the endpoints of the tunnels you want to handle.	Network conditions constrain the source and destination hosts in each connection.
Direction	Bidirectional or unidirectional (configurable). Tunnel rules are bidirectional by default, so they can handle all traffic between tunnel endpoints.	Unidirectional only (nonconfigurable). Prefilter rules match source-to-destination traffic only. Return traffic for allowed connections is also permitted.
Rezone sessions for further analysis	Supported, using tunnel zones; see Using tunnel zones to apply access control at the tunnel level, on page 2534.	Not applicable.

How the system processes plain-text tunnels

Some protocols, such as GRE, create plain-text (not encrypted) tunnels that contain (or encapsulate) multiple inner connections. These connections might flow between discontinuous networks. These tunnels are especially useful for routing custom protocols over IP networks, IPv6 traffic over IPv4 networks, and so on.

Summary

A plain-text tunnel has two sets of IP headers:

- Outer headers: Specify tunnel endpoint addresses for prefilter policy handling
- Inner headers: Contain actual connection endpoint addresses for granular inspection

Workflow

The system processes plain-text tunnels through these stages:

1. The system receives a plain-text tunnel with dual IP header structure.
   - Outer header—An outer encapsulation header specifies the source and destination IP addresses of the tunnel endpoints—the routed interfaces of the network devices on either side of the tunnel.
   - Inner headers—Inner payload headers specify the source and destination IP addresses of the encapsulated connections' actual endpoints. There are separate headers for each connection contained within the tunnel.

   For regular, non-tunneled connections, there is only the "outer" header, as the connection does not contain any subordinate connections.
2. The prefilter policy processes the outer tunnel headers to determine handling approach.

   The prefilter policy uses outer headers to handle traffic. You can block or fast path entire plaintext, passthrough tunnels at this stage.

   If you do not want to allow plain-text tunnels at all, you can set the default action in the prefilter policy to block all plain-text tunnels. This block does not apply to encrypted tunnels such as VPN.
3. The system exposes inner headers for tunnels that pass prefilter analysis.

   Any tunnels that you analyze (that is, not block or fast path) have these inner headers exposed for use by the rest of the security policies and other features such as QoS. These features use the innermost detectable level of headers to ensure the most granular level of inspection and handling possible. This ensures that bad actors cannot hide illegitimate traffic within a plain-text tunnel.

   Alternatively, you can use tunnel zones to re-tag tunneled traffic so that you can write access control rules specifically for select tunnels. This lets you customize intrusion inspection or other layer 7 filtering to connections contained with a tunnel. See Using tunnel zones to apply access control at the tunnel level, on page 2534.

Using tunnel zones to apply access control at the tunnel level

A tunnel zone represents certain types of plaintext, passthrough tunnels that you explicitly tag for special analysis. Although a tunnel zone is not an interface object, you can use it as an interface constraint in some configurations.

Tunnel zone implementation

As explained in How the system processes plain-text tunnels, on page 2533, once a plain-text tunneled connection passes through the prefilter policy, the access control policy applies its rules to each connection contained within the tunnel.

However, you might want to apply the same policy to all the connections within a tunnel. For example, you might have a special intrusion policy you want to use for tunneled connections.

To make it possible to apply access control rules to all connections within a tunnel, you must use the prefilter policy to re-zone the tunnel. Rezoning applies a tunnel zone tag to the plain-text tunnel. You can then use that tunnel zone instead of a security zone when you define the source interface criterion of an access control rule.

Because a tunnel zone does not contain interfaces, the source interface through which the tunnel enters the device is not relevant. You can apply a tunnel zone tag to multiple tunnels that enter the device through different interfaces and access control treats them all in the same way.

Note Connections in rezoned tunnels do not match security zone constraints in an access control rule. After you rezone a tunnel, access control rules match encapsulated connections with their newly assigned tunnel zone, not the original security zone.

See Rezone tunnels for customized inspection, on page 2552 for a walkthrough of a tunnel zone implementation, and a discussion of the implications of rezoning without explicitly handling rezoned traffic.

Large flow offloads

A large flow offload is a hardware-based traffic processing feature that

- handles certain fastpathed traffic in the NIC hardware rather than in the Firewall Threat Defense software
- results in higher throughput and lower latency for data-intensive applications such as large file transfers, and
- includes both static flow offload (based on prefilter policy configuration) and dynamic flow offload (based on trust and other criteria).

Supported models and offload types

For the following models, certain traffic that you configure to be fastpathed by a prefilter policy is handled by the hardware (specifically, in the NIC), not by the Firewall Threat Defense software:

- Firepower 4100/9300
- Secure Firewall 3100 (static only)
- Secure Firewall 4200 (static only)

Offloaded flows continue to receive limited stateful inspection, such as basic TCP flag and option checking. The system can selectively escalate packets to the firewall system for further processing if necessary.

Two types of flow offload are available:

- *Static flow offload*: Offloads flows based on prefilter policy configuration that applies the Fastpath action.
- *Dynamic flow offload*: By default, devices offload flows based on other criteria, including trust.

Applications that can benefit from offloading large flows include:

- High Performance Computing (HPC) Research sites, where the firewall is deployed between storage and high compute stations. When one research site backs up using FTP file transfer or file sync over NFS,

the large amount of data traffic affects all connections. Offloading FTP file transfer and file sync over NFS reduces the impact on other traffic.

- High Frequency Trading (HFT), where the firewall is deployed between workstations and the Exchange, mainly for compliance purposes. Security is usually not a concern, but latency is a major concern.

Flows that can be offloaded include:

- (Static flow offload only.) Connections that are fastpathed by the prefilter policy.
- Standard or 802.1Q tagged Ethernet frames only.
- Cluster asymmetric flows (Secure Firewall 4200). For asymmetric flows, cluster redirect lets the forwarding node offload flows to hardware. When traffic for an existing flow is sent to a different node, then that traffic is redirected to the owner node over the cluster control link. Because asymmetric flows can create a lot of traffic on the cluster control link, letting the forwarder offload these flows can improve performance.
- (Dynamic flow offload only):
  - Inspected flows that the inspection engine decides no longer need inspection. These flows include:
    - Flows handled by access control rules that apply the Trust action and are based on security zone, source and destination network and port matching only.
    - TLS/SSL flows that are not selected for decryption using a decryption policy.
    - Flows that are trusted by the Intelligent Application Bypass (IAB) policy either explicitly or due to exceeding flow bypass thresholds.
    - Flows that match file or intrusion policies that result in trusting the flow.
    - Any allowed flow that no longer needs to be inspected.
  - IPS preprocessor inspected flows:
    - SSH and SMTP.
    - FTP preprocessor secondary connections.
    - Session Initiation Protocol (SIP) preprocessor secondary connections.
  - Intrusion rules that use keywords (also referred to as *options*)



Important For details, exceptions, and limitations to the above, see Restrictions for flow offload implementation, on page 2537.

Static flow offload configuration

To offload eligible traffic to hardware, create a prefilter policy rule that applies the **Fastpath** action. Use prefilter rules for TCP/UDP, and tunnel rules for GRE.

(Not recommended.) To disable flow offload, use FlexConfig:

- Disable static (and thereby dynamic) flow offload: **no flow-offload enable**

- Re-enable flow offload: **flow-offload enable**
- Disable cluster redirection: **no flow-offload cluster-redirect**

For more information, see *Cisco Secure Firewall ASA Series Command Reference*.

Dynamic flow offload configuration

Dynamic flow offload is enabled by default. However, dynamic offload occurs only if static flow offload is enabled, regardless of whether prefiltering is configured.

To disable dynamic flow offload, use the device CLI:

- Disable dynamic offload: **configure flow-offload dynamic whitelist disable**
- Re-enable dynamic offload: **configure flow-offload dynamic whitelist enable**

Restrictions for flow offload implementation

Device limitations

The feature is supported on the following devices:

- Secure Firewall 3100 (static only)
- Secure Firewall 4200 (static only)
- Firepower 4100/9300

Flows that cannot be offloaded

Do not attempt to offload these flow types:

- Any flows that do not use IPv4 addressing, such as IPv6 addressing.
- Flows for any protocol other than TCP, UDP, and GRE.

Note PPTP GRE connections cannot be offloaded.

- Flows on interfaces configured in passive, inline, or inline tap mode. Routed and switched interfaces are the only types supported.
- (3100/4200.) Offload based on inner header for tunnelled flows.
- (3100/4200.) Multi-instance offload.
- Flows that require inspection by Snort or other inspection engines. In some cases, such as FTP, the secondary data channel can be offloaded although the control channel cannot be offloaded.
- IPsec and TLS/DTLS VPN connections that terminate on the device.

Note The Secure Firewall 3100 and Secure Firewall 4200 platforms are exceptions to this restriction. These platforms support flow-offload for IPsec and DTLS VPN connections terminating to the device. For detailed information on this feature and its implementation, please refer to the release notes specific to the Secure Firewall 3100 and Secure Firewall 4200 platforms.

- Flows that require encryption or decryption. For example, connections decrypted due to a decryption policy.
- Multicast flows in routed mode. They are supported in transparent mode if there are only two member interfaces in a bridge group.
- TCP Intercept flows.
- TCP state bypass flows. You cannot configure flow offload and TCP state bypass on the same traffic.
- Flows tagged with security groups.
- Reverse flows that are forwarded from a different cluster node, in the case of asymmetric flows in a cluster.
- Centralized flows in a cluster, if the flow owner is not the control unit.
- Flows that include IP options cannot be dynamically offloaded.

Additional limitations

Consider these operational restrictions when implementing flow offload:

- Flow offload and Dead Connection Detection (DCD) are not compatible. Do not configure DCD on connections that can be offloaded.
- If more than one flow that matches flow offload conditions are queued to be offloaded at the same time to the same location on the hardware, only the first flow is offloaded. The other flows are processed normally. This is called a *collision*. Use the **show flow-offload flow** command in the CLI to display statistics for this situation.
- Dynamic flow offload disables all TCP normalizer checks.
- Although offloaded flows pass through FXOS interfaces, statistics for these flows do not appear on the logical device interface. Thus, logical device interface counters and packet rates do not reflect offloaded flows.

Cluster redirection

When traffic for an existing flow is sent to a different node, then that traffic is redirected to the owner node over the cluster control link. Because asymmetric flows can create a lot of traffic on the cluster control link, offloading these flows can improve performance, for example, for distribited site-to-site VPN mode. This feature is supported for these environments:

- Secure Firewall 4200
- Routed and transparent mode

- Spanned EtherChannel mode

Conditions for reversing offload

After a flow is offloaded, packets within the flow are returned to the firewall operating system for further processing if they meet these conditions:

- They include TCP options other than Timestamp.
- They are fragmented.
- They are subject to Equal-Cost Multi-Path (ECMP) routing, and ingress packets move from one interface to another.
- A change to the routing table applies to the flow. A packet is returned to the device to determine the new route.

Prefilter vs access control policy

Prefilter and access control policies both allow you to block and trust traffic. In a prefilter policy, the "trust" functionality is called "fastpathing" because it skips more inspection. If you do not configure a custom prefilter policy, you can only approximate, but not replicate, prefilter functionality using Block and Trust rules placed early in the access control policy.

Characteristic	Prefilter Policy	Access Control Policy
Primary function	Quickly fastpath or block certain types of plaintext or passthrough tunnels, or tailor subsequent inspection to encapsulated traffic. Fastpath or block any other connections that benefit from early handling.	Inspect and control all network traffic, using simple or complex criteria, including contextual information and deep inspection results.
Sequence within access control	First The system matches traffic to prefilter criteria before all other access control configurations. Thus, fast-pathing in a prefilter policy skips many other policies and inspections.	Last Other policies that might inspect a connection, including Security Intelligence and decryption, are applied before access control rules are evaluated.
Rule actions	Fewer You can stop further inspection (Fastpath and Block) or allow further analysis with the rest of access control (Analyze).	More Access control rules have a larger variety of actions, including monitoring, deep inspection, block with reset, and interactive blocking. See Access control rule actions, on page 2046.

Characteristic	Prefilter Policy	Access Control Policy
Bypass capability	Fastpath rule action Fastpathing traffic in the prefilter stage bypasses all further inspection and handling. This includes: • Security Intelligence • User identity discovery (identity policy) • Decryption • Access control rules • Deep inspection of packet payloads • Discovery • Rate limiting	Trust rule action Traffic trusted by access control rules is only exempt from intrusion inspection and discovery. If another policy that requires inspection is applied to a connection that is eventually trusted, the connection is inspected anyway. Trust means that intrusion and file policies are not applied to the connection.
Rule criteria	Limited Rules in the prefilter policy use simple network criteria: IP address, VLAN tag, port, and protocol. For tunnels, tunnel endpoint conditions specify the IP address of the routed interfaces of the network devices on either side of the tunnel.	Robust Access control rules use network criteria, but also user, application, requested URL, and other contextual information available in packet payloads. Network conditions specify the IP address of source and destination hosts.
IP headers used (tunnel handling)	Outermost Using outer headers allows you to handle entire plaintext, passthrough tunnels. For nonencapsulated traffic, prefiltering still uses "outer" headers—which in this case are the only headers.	Innermost possible For a nonencrypted tunnel, access control acts on its individual encapsulated connections, not the tunnel as a whole. See How the system processes plain-text tunnels, on page 2533.
Rezone encapsulated connections for further analysis	Rezones tunneled traffic Tunnel zones allow you to tailor subsequent inspection to prefiltered, encapsulated traffic.	Uses tunnel zones Access control uses the tunnel zones you assign during prefiltering. See Using tunnel zones to apply access control at the tunnel level, on page 2534.
Connection logging	Fastpathed and blocked traffic only. Allowed connections may still be logged by other configurations.	Any connection

Requirements and prerequisites for prefilter policies

This reference provides the requirements and prerequisites necessary for configuring prefilter policies, including supported models, domain requirements, and required user roles.

Model support

Firewall Threat Defense

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Guidelines and limitations for the prefilter policy

You can log connections fast pathed and blocked by the prefilter policy. Prefilter connection events contain information on whether and how logged connections—including entire tunnels—were prefiltered. However, because fast pathed and blocked connections are not subject to deep inspection, associated connection events contain limited information.

Access control uses a hierarchical implementation that complements multitenancy. Along with other advanced settings, you can lock a prefilter policy association, enforcing that association in all descendant access control policies. For more information, see Access control policy inheritance, on page 2000.

When you deploy a prefilter policy, its rules are not applied on the existing connections. Hence, traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic on an existing connection that could match a policy is omitted from the hit count. To have the rules effectively applied, clear the existing connections, and then deploy the policy.

For a prefilter rule with action as Fastpath and logging options disabled, you might observe that the end-of-flow events are still generated in the system. However, the events are not visible on the management center event pages.

Guidelines and limitations for encapsulated traffic handling

Supported encapsulated traffic types

This topic discusses guidelines for the following types of encapsulated traffic:

- Generic Routing Encapsulation (GRE)
- Point-to-Point Protocol (PPTP). (This protocol is not handled by tunnel rules in the prefilter policy.)
- IPinIP
- IPv6inIP
- Teredo

GRE tunnel limitations

GRE tunnel processing is limited to IPv4 and IPv6 passenger flows. Other protocols, such as PPTP and WCCP, are not supported within the GRE tunnel.

GRE v1 and PPTP bypass restrictions

GRE v1 (sometimes referred to as *stateful GRE*) and PPTP traffic bypass outer flow processing.

- Sessions are over a single tunnel that is not load-balanced.
- There is no HA or clustering replication.
- Primary and secondary flow relationships are not maintained.
- Prefilter policy white and black lists are not supported.

Passenger flow processing is supported for IPv6inIP and Teredo but the following limitations apply:

GRE v0 sequence number field requirement

All endpoints sending traffic on the network must send GREv0 traffic with the sequence number field as optional; otherwise, the sequence number field is removed. RFC 1701 and RFC 2784 both specify the sequence field as optional.

Tunnel interface compatibility

Prefilter and access control policy rules are applied to all tunnel types on routed, transparent, inline-set, inline-tap, and passive interfaces.

Protocol references

For more information about the GRE and PPTP protocols, see the following:

- RFC 1701, RFC 2784, and RFC 2890 (GRE protocol v0)
- RFC 2637 (PPTP and GRE protocol v1)

Best practices for fastpath prefiltering

When you use the fastpath action in a prefilter rule, the matching traffic bypasses inspection and is simply transmitted through the device. Use this action for traffic that you can trust and that would not benefit from any of the security features available.

- VPN traffic that is going through the device. That is, the device is not an endpoint in the VPN topology.
- SQL traffic between trusted endpoints on the internal network.
- Scanner traffic. Scanner probes can create a lot of false-positive responses from intrusion policies.
- Voice/video.
- Backups.
- Management traffic (sftunnel) that traverses Firewall Threat Defense devices. Performing deep inspection on management traffic (using access control policies) can cause issues. You can prefilter based on port TCP/8305 between the management center and managed devices.

Configure prefilter policies

Configure prefilter policies to perform custom prefiltering and to control tunnel traffic analysis on managed devices.

To perform custom prefiltering, configure prefilter policies and assign them to access control policies. It is through the access control policy that prefilter policies get assigned to managed devices.

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Prefilter**.

At the top of the page, there are convenient links to some related features: object management and access control.

You cannot tell from the prefilter page, or when editing a prefilter policy, whether it is used by an access control policy or which policies use it.

Step 2 Do any of the following:

- Create—Click **New Policy** to create a custom prefilter policy. You are prompted for a name and an optional description. The new policy is opened for edit after you click **Save**. A new prefilter policy has no rules and a default action of Analyze all tunnel traffic.
- Edit—Click **Edit** ( ) to open a policy so that you can add rules or otherwise modify it.
- Copy—Click **Copy** ( ). You are prompted for a name. The copy is added to the list when you click **OK**, but it is not opened for edit.
- Delete—Click **Delete** ( ) to remove a policy. You cannot delete a policy that is currently being used by an access control policy.

Step 3 When editing a prefilter policy, you can do the following:

- Set the default action—Set the default action for tunnels at the bottom of the policy. See Configuring the default action, on page 2545.
- Add prefilter rules—Click **Add Prefilter Rule**. The rule is added to the bottom of the policy. See Configure prefilter rules, on page 2545.
- Add tunnel rules—Click **Add Tunnel Rule**. The rule is added to the bottom of the policy. See Configure tunnel rules, on page 2548.
- Insert rules—Right click the rule before where you want to add the new rule, and select either **Insert New Prefilter Rule** or **Insert New Tunnel Rule**, as appropriate.

- Copy/cut and paste rules—To copy a rule, right-click the rule and select **Copy**. You can also cut the rule (remove it) by selecting **Cut**. Then, select where you want to insert the rule, right click, and select either **Paste Above** or **Paste Below**. When copying, a number in parentheses is added to the rule name; you should edit the rule and change the name and other criteria as needed. Cut rules retain the original name.
- Edit a rule—Click **Edit** ( ) to open the rule editor. You can also click on any element for the rule in the table. See Configure prefilter rules, on page 2545 and Configure tunnel rules, on page 2548.
- Delete rules—Click **Delete** ( ) to remove a rule.
- Move a rule—Click a rule and drag and drop it to the right location. Alternatively, cut the rule and then paste it in the right location. You can also change the rule location while editing it.

  Properly creating and ordering rules is a complex task, but one that is essential to building an effective deployment. If you do not plan carefully, rules can preempt other rules or contain invalid configurations. For more information, see Prefilter policy rule order, on page 2532.
- Move a prefilter rule to an associated access control policy—To move a prefilter rule from the prefilter policy to an associated access control policy, right-click the rule and select **Move to another policy**. You cannot move tunnel rules to an access control policy. See Move prefilter rules to an access control policy, on page 2551.
- View hit counts—To view statistics on how many connections matched each rule, click **Analyze Hit Counts**. See View rule hit counts, on page 2018.
- Search for rules—Use the **Search Rules** box to do a plain text search. Type the string and press enter. The search box shows the number of hits and up/down buttons to move through them.

  Matches are highlighted in the table. Matches include partial hits, but only to explicitly set items. For example, entering "any" does not match the default network, but "any-ipv4" will match because it is a network object and not a passive default. Action and rule type do not match the search string.
- View object details—To see the contents of an object used in a rule, right click the object and select **Object Details**.
- Change the policy description—To modify the description of a policy, click in the description text below the policy name at the top of the page. Click away from the description to save changes.

Be sure to click **Save** to save any changes you make.

Step 4 Use the prefilter policy in the appropriate access control policies.

For each access control policy that should use a custom prefilter policy:

- Edit the access control policy and select the prefilter policy in the **Prefilter Policy** link or in the advanced settings. See Associate prefilter, decryption, and identity policies with an access control, on page 2016.
- If you applied tunnel zone tags in any tunnel rules, create access control rules that use those tunnel zones as source interface criteria.

What to do next

Deploy the changes.

If you deploy time-based rules, specify the time zone of the device to which the policy is assigned. See Configure a time zone for a device, on page 1304.

Configuring the default action

Use the default action to determine how plain-text encapsulated tunnels that do not match any tunnel rules should be handled.

The default action does not apply to non-tunnel traffic. Any other connections that do not match prefilter rules are passed on for processing by subsequent policies.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Prefilter** and create or edit a prefilter policy.

Step 2 Choose a **Default Action: Tunneled Traffic** option:

a. **Analyze all tunnel traffic**—These tunnels are allowed and their encapsulated connections are evaluated by the access control and other policies

Block all tunnel traffic—Tunnels are dropped and no longer evaluated.

Step 3 If you select **Block all tunnel traffic**, click **Logging** ( ) next to the default action.

You can configure default action logging for blocked tunnels only.

Configure prefilter rules

Configure prefilter rules to provide early blocking for unwanted traffic and to fast path traffic that does not benefit from access control inspection.

Use prefilter rules in a non-default prefilter policy to provide early blocking for unwanted traffic, and to fast path traffic that does not benefit from access control inspection.

If your device supports flow offload, fast path traffic is eligible for flow offload if it meets all other criteria. See Large flow offloads, on page 2535.

Before you begin

You cannot create rules in the default prefilter policy. If you have not created a custom prefilter policy, do so now.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Prefilter** and create or edit a prefilter policy.

Click **Add Prefilter Rule**.

You can also right-click a rule and select **Insert New Prefilter Rule**.

Select whether the rule should be **Enabled**. A rule must be enabled to affect traffic through the device.

Step 2 Select the **Action** to be taken on matching traffic.

- Fastpath—Exempts matching traffic from all further inspection AND control, including access control, identity requirements, AND rate limiting. These connections are eligible for flow offload if the device supports it.
- Block—Blocks matching traffic without further inspection of any kind.
- Analyze—Allows traffic to continue to be analyzed by the rest of access control. If passed by access control AND any related deep inspection, this traffic may also be rate limited.

Step 3 Select where to **Insert** the rule.

The default is to insert the rule at the end of the policy, but you can select an existing rule and insert the new rule before or after that rule. The action of the first rule that matches a connection is applied to that connection, so order matters.

Step 4 (Optional.) Select the object that defines the **Time Range** for the rule.

You can limit what times of day, or which days, a rule is operational. For example, the rule could apply during business hours only. Select an existing time range object OR create a new one. For more information, see Create a time range object, on page 3036.

Step 5 Configure the traffic matching conditions for the rule.

The rule is applied to connections that match all of the conditions. Within a condition type, separate elements are OR'ed, so that matching any of the selected items counts as a match. For the rule as a whole, all condition types are AND'ed: a connection must match at least one item on each tab to match the rule. You can use these types of conditions:

- **Interface Objects**—Select the security zones OR interface groups that contain the interfaces to which the rule should apply. The default, no objects selected, applies the rule to all interfaces on the device.
- **Networks**—Select the network objects that define the source (**Add to Source**) OR destination (**Add to Destination**) IP addresses, OR both. The default (no objects selected) matches all source AND destination IP addresses.
- **VLAN tags**—Select the objects that define the VLANs to which the rule should apply. For more information, see VLAN tags rule conditions, on page 2547.
- **Ports**—Select the TCP/UDP ports, OR other protocol, for the connections AND add them to the source OR destination criteria. Source criteria is TCP/UDP only, but you can select other protocols for the destination. For more information, see Port rule conditions for prefilter rules, on page 2547.

Step 6 (Fastpath AND Block rules only.) Click **Logging** and configure how matching connections should be logged.

Step 7 (Optional.) Click **Comment** and add comments to the rule. Use comments to help you understand the purpose of the rule and its change history.

Step 8 Click **Add**. If you need to move the rule, do so now.

Step 9 Click **Save** to save the policy.

VLAN tags rule conditions

VLAN tags rule conditions are filtering mechanisms to control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic.

VLAN tag filtering behavior

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note this Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

Port rule conditions for prefilter rules

Port conditions match traffic based on the source and destination ports.

Port condition types

Depending on the rule type, "port" can represent any of these types:

- TCP and UDP—You can control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22. TCP and UDP are the only source ports supported. You can also specify them as destination ports.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3. When you select ICMP as a destination port, you are prompted to specify the type and code.
- Protocol—You can control traffic using other protocols that do not use ports. Not all protocols are supported. You must select the protocol as a destination port.

Source and destination port constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. However, when matching GRE, IP in IP, IPv6 in IP, or Teredo Port 3544, configure tunnel rules instead of port conditions in prefilter rules.

Configure tunnel rules

Configure tunnel rules to control how encapsulated tunnel traffic is processed by the firewall system.

- Provide early blocking for unwanted encapsulated tunnel traffic.
- Fast path tunnels that do not benefit from access control inspection.
- Place encapsulated tunnels into a tunnel zone, called rezoning, so that you can write access control rules to apply inspection to the connections within the tunnel as a whole.

The prefilter policy's default action applies to encapsulated tunnels. Thus, create rules for tunnel traffic only to apply different actions to specific tunnels.

For example, if the default action is to analyze tunnels, this means that all tunnels are passed on to the access control policies AND other policies that inspect, allow, AND potentially block connections, AND these policies will evaluate the inner connections. If that default action is what you want for all tunnels, then you do not need to create any tunnel rules.

If instead:

- The default action is to block encapsulated tunnels, use rules to exempt specific tunnels from blocking AND either pass them to other access control policies (analyze action) OR bypass any additional processing (fastpath).
- There are tunnels where you want to apply access control rules to all encapsulated connections, create rules to rezone the tunnels to the appropriate tunnel zone tag AND also write rules in the access control policy for those tunnel zone tags.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Prefilter**. Create or edit a prefilter policy.

Step 2 Click **Add Tunnel Rule**.

You can also right-click a rule AND select **Insert New Tunnel Rule**.

a) Enter a **Name** for the rule.

b) Select whether the rule should be **Enabled**.

A rule must be enabled to affect traffic through the device.

c) Select where to **Insert** the rule.

The default is to insert the rule at the end of the policy, but you can select an existing rule AND insert the new rule before OR after that rule. The action of the first rule that matches a connection is applied to that connection, so order matters.

d) (Optional.) Select the object that defines the **Time Range** for the rule.

You can limit what times of day, OR which days, a rule is operational. For example, the rule could apply during business hours only. Select an existing time range object OR create a new one. For more information, see Create a time range object, on page 3036.

Step 3 Select the **Action** to be taken on matching traffic.

- Fastpath—Exempts matching traffic from all further inspection AND control, including access control, identity requirements, AND rate limiting. Fastpathing a tunnel fastpaths all encapsulated connections. These connections are eligible for flow offload if the device supports it.
- Block—Blocks matching traffic without further inspection of any kind. Blocking a tunnel blocks all encapsulated connections.
- Analyze—Allows traffic to continue to be analyzed by the rest of access control using inner headers.. If passed by access control AND any related deep inspection, this traffic may also be rate limited.

a) (Analyze rules only; when rezoning only.) Select a tunnel zone tag in **Assign Tunnel Zone** if your intention is to rezone the tunnel.

Leave this option blank if you want access control to apply to the encapsulated connections within the tunnel without special processing. For more information about rezoning AND tunnel group tags, see:

- Using tunnel zones to apply access control at the tunnel level, on page 2534
- Creating tunnel zones, on page 2550
- Rezone tunnels for customized inspection, on page 2552

Caution
Exercise caution when assigning tunnel zones. Connections in rezoned tunnels may not match security zone constraints in later evaluation. See Rezone tunnels for customized inspection, on page 2552 for a brief walkthrough of a tunnel zone implementation, AND a discussion of the implications of rezoning without explicitly handling rezoned traffic.

Step 4 Select one of these options to set the direction for the rule:

- **Match tunnels only from source** (unidirectional)—Match source-to-destination traffic only. Matching traffic must originate from one of the specified source interfaces OR tunnel endpoints, AND leave through one of the destination interfaces OR tunnel endpoints. Return traffic for allowed connections is also permitted. (Note that prefilter rules are always unidirectional.)
- **Match tunnels from source AND destination** (bidirectional)—Match both source-to-destination traffic AND destination-to-source traffic. The effect is identical to writing two unidirectional rules, one the mirror of the other. This is the default.

Step 5 Click the **Encapsulation & Ports** tab and select the protocols for this rule.

You must select at least one of these options in a tunnel rule to specify the encapsulation protocol used. The list includes the IP protocol number for the option.

- **GRE** (47)
- **IP-in-IP** (4)
- **IPv6-in-IP** (41)
- **Teredo** (UDP (17)/3455)

Step 6 Configure the traffic matching conditions for the rule.

The rule is applied to connections that match all of the conditions. Within a condition type, separate elements are OR'ed, so that matching any of the selected items counts as a match. For the rule as a whole, all condition types are AND'ed: a connection must match at least one item on each tab to match the rule. You can use these types of conditions:

- **Interface Objects**—Select the security zones OR interface groups that contain the interfaces to which the rule should apply. The default, no objects selected, applies the rule to all interfaces on the device.
- **Tunnel Endpoints**—The network objects that define the source (**Add to Source**) AND destination (**Add to Destination**) IP addresses of the tunnel. The default (no objects selected) matches all source AND destination IP addresses.
- **VLAN tags**—Select the objects that define the VLANs to which the rule should apply. For more information, see VLAN tags rule conditions, on page 2547.

Step 7 (Fastpath AND Block rules only.) Click **Logging** AND configure how matching connections should be logged.

Step 8 (Optional.) Click **Comment** AND add comments to the rule.

Use comments to help you understand the purpose of the rule AND its change history.

Step 9 Click **Add**.

If you need to move the rule, do so now.

Click **Save** to save the policy.

Creating tunnel zones

Create a tunnel zone tag to represent certain types of plaintext, passthrough tunnels that you explicitly tag for special analysis.

A *tunnel zone* represents certain types of plaintext, passthrough tunnels that you explicitly tag for special analysis. A tunnel zone is not an interface object, even though you can use it as an interface constraint in access control rules, as explained in Using tunnel zones to apply access control at the tunnel level, on page 2534.

You can create zones in the object manager or when editing a tunnel rule.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Tunnel Zone** from the list of object types.

Step 3 Click **Add Tunnel Zone**.

You can also edit an existing tunnel or delete a tunnel zone that is not being used in a tunnel rule.

Step 4 Enter a **Name** and, optionally, a **Description**.

Step 5 Click **Save**.

What to do next

- Assign the tunnel zone to plaintext, passthrough tunnels as part of a tunnel rule. See Configure tunnel rules, on page 2548.
- Use the tunnel zone as source interface criteria in an access control rule.

Move prefilter rules to an access control policy

Moving prefilter rules to an access control policy allows you to reposition rules for better policy organization and effectiveness when the rules would function more appropriately within the access control policy structure.

You can move prefilter rules from a prefilter policy to the associated access control policy.

Before you begin

Note the following conditions before you proceed:

- Only prefilter rules can be moved to an access control policy. Tunnel rules cannot be moved.
- The prefilter rules can be moved only to an access control policy that uses the prefilter policy. Because you can use the same prefilter policy in multiple access control policies, you need to select the policies to which the rule should be moved.
- The prefilter rules that use interface groups to define interface criteria cannot be moved. You can move only those rules that use security zones, or that include no interface criteria.
- The **Action** parameter in the prefilter rule is changed to a suitable action in the access control rule when moved.
  - **Analyze** becomes allow.
  - **Block** becomes block.
  - **Fastpath** becomes trust.
- The logging configuration in the rule is retained.
- The comments in the prefilter rule configuration are lost after moving the rule. However, a new comment is added in the moved rule mentioning the source prefilter policy.

- While moving rules from the source policy, if another user modifies those rules, you are warned. You can continue with the process after refreshing the page.

Procedure

Step 1 In the prefilter policy editor, select the rules that you want to move with a left-click on your mouse.

Tip
To select multiple rules, use Shift+click.

Step 2 Right-click the selected rules and choose **Move to another policy**.

Step 3 Select the destination access control policies and click **Add** to add them to the list of selected policies.

Step 4 From the **Place Rules** drop-down list, choose where you want to position the moved rules:

- To position as the last set of rules in the **Default** section, choose **At the bottom (within the Default section)**.
- To position as the first set of rules in the **Mandatory** section, choose **At the top (within the Mandatory section)**.

Step 5 Click **Move**.

What to do next

- Deploy configuration changes.

Examples of prefilter policy

These topics provide examples of using the prefilter policy.

Rezone tunnels for customized inspection

Rezone GRE tunnels to enable customized inspection policies for encapsulated traffic while maintaining different inspection policies for other traffic in the same security zone.

This example procedure summarizes how you might rezone GRE tunnels for further analysis, using tunnel zones. You can adapt the concepts described here to other scenarios requiring tailored traffic inspection of connections encapsulated in plaintext or passthrough tunnels.

Consider a situation where your organization's internal traffic flows through the Trusted security zone. The Trusted security zone represents a set of interfaces across multiple managed devices deployed in various locations. Your organization's security policy requires that you allow internal traffic after deep inspection for exploits and malware.

Internal traffic sometimes includes plaintext, passthrough, GRE tunnels between particular endpoints. Because the traffic profile of this encapsulated traffic is different from your "normal" interoffice activity—perhaps it is known and benign—you can limit inspection of certain encapsulated connections while still complying with your security policy.

In this example, after you deploy configuration changes:

- Plaintext, passthrough, GRE-encapsulated tunnels detected in the Trusted zone have their individual encapsulated connections evaluated by one set of intrusion and file policies.
- All other traffic in the Trusted zone is evaluated with a different set of intrusion and file policies.

You accomplish this task by *rezoning* GRE tunnels. Rezoning ensures that access control associates GRE-encapsulated connections with a custom *tunnel* zone, rather than their original Trusted *security* zone. Rezoning is required due to the way access control handles encapsulated traffic; see How the system processes plain-text tunnels, on page 2533 and Using tunnel zones to apply access control at the tunnel level, on page 2534.

Procedure

Step 1 Configure custom intrusion and file policies that tailor deep inspection to encapsulated traffic, and another set of intrusion and file policies tailored to nonencapsulated traffic.

Step 2 Configure custom prefiltering to rezone GRE tunnels flowing through the Trusted security zone.

Create a custom prefilter policy and associate it with access control. In that custom prefilter policy, create a tunnel rule (in this example, **`GRE_tunnel_rezone`**) and a corresponding tunnel zone (**`GRE_tunnel`**). For more information, see Configure tunnel rules, on page 2548.

Table 134: GRE_tunnel_rezone tunnel rule

Rule Component	Description
Interface object condition	Match internal-only tunnels by using the Trusted security zone as both a Source Interface Object and Destination Interface Object constraint.
Tunnel endpoint condition	Specify the source and destination endpoints for the GRE tunnels used in your organization. Tunnel rules are bidirectional by default. If you do not change the **Match tunnels from...** option, it does not matter which endpoints you specify as source and which as destination.
Encapsulation condition	Match GRE traffic.
Assign Tunnel Zone	Create the **`GRE_tunnel`** tunnel zone, and assign it to tunnels that match the rule.
Action	Analyze (with the rest of access control).

Step 3 Configure access control to handle connections in rezoned tunnels.

In the access control policy deployed to your managed devices, ensure that the prefilter policy is used. Then, configure a rule (in this example, **`GRE_inspection`**) that handles the traffic you rezoned. For more information, see Create and edit access control rules, on page 2066.

Table 135: GRE_inspection access control rule

Rule Component	Description
Security zone condition	Match rezoned tunnels by using the GRE_tunnel tunnel zone as a Source Zone constraint.
Action	Allow, with deep inspection enabled. Choose the file and intrusion policies tailored to inspect encapsulated internal traffic.

Caution
If you skip this step, the rezoned connections may match **any** access control rule not constrained by security zone. If the rezoned connections do not match any access control rules, they are handled by the access control policy default action. Make sure this is your intent.

Step 4 Configure access control to handle nonencapsulated connections flowing through the Trusted security zone.

In the same access control policy, configure a rule (in this example, **`internal_default_inspection`**) that handles non-rezoned traffic in the Trusted security zone.

Table 136: Internal_default_inspection access control rule

Rule Component	Description
Security zone condition	Match non-rezoned internal-only traffic by using the Trusted security zone as both a Source Zone and Destination Zone constraint.
Action	Allow, with deep inspection enabled. Choose the file and intrusion policies tailored to inspect nonencapsulated internal traffic.

Step 5 Evaluate the position of the new access control rules relative to preexisting rules. Change rule order if necessary.

If you place the two new access control rules next to each other, it does not matter which you place first. Because you rezoned GRE tunnels, the two rules cannot preempt each other.

Step 6 Save all changes.

What to do next

- Deploy configuration changes.

Offload large flows

This task helps improve performance quickly for large connections that you know you can trust by proactively identifying connections as being offload eligible through matching them to a prefilter fastpath rule.

You can use the prefilter policy to configure static flow offload. If your hardware model supports it, flow offload moves eligible fast path connections to NIC processing, avoiding performance intensive inspection.

Some hardware models will dynamically identify off-loadable connections automatically, but you can proactively identify connections as being offload eligible by matching them to a prefilter fastpath rule. This static configuration helps you improve performance quickly for large connections that you know you can trust.

In the following example, it is assumed that two servers are behind sufficient protection so that you can trust the connections between them. Perhaps one server is using the other for backups. Thus, you can write a fastpath rule for any traffic going from the main server, 10.100.1.20/24, to the backup server, 10.100.2.2/24. The example assumes the servers are accessible through regular, physical routed interfaces.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Prefilter** and create or edit a custom prefilter policy.

Step 2 Click **Add Prefilter Rule**.

You can also right-click a rule and select **Insert New Prefilter Rule**.

Step 3 Enter a **Name** for the rule. For example, Offload_Server_A.

Step 4 In **Action**, select **Fastpath**.

Step 5 In **Insert**, select the rule above or below which the rule should be inserted.

You can also move the rule after you create it.

Step 6 Click the **Networks** tab and specify the source and destination servers.

- In the edit box below **Source Networks**, type 10.100.1.20 and click **Add**.
- In the edit box below **Destination Networks**, type 10.100.2.2 and click **Add**.

Step 7 (Optional.) Configure the other traffic matching characteristics for the flow.

Click these tabs to fill in the criteria. Only the Networks tab needs to be configured for this example, so configure these settings only if they make sense for your situation.

- **Interface objects**—The security zones or interface objects that define the source or destination interfaces, or both. This specification is not needed for this example.
- **VLANs**—Because this example uses routed physical interfaces instead of switch interfaces, this setting is not relevant. Even in those circumstances, you do not need to limit the rule based on VLAN.
- **Ports**—If you want to limit the offload to specific TCP/UDP ports, you can specify them. This limits the scope of the rule to these ports only.

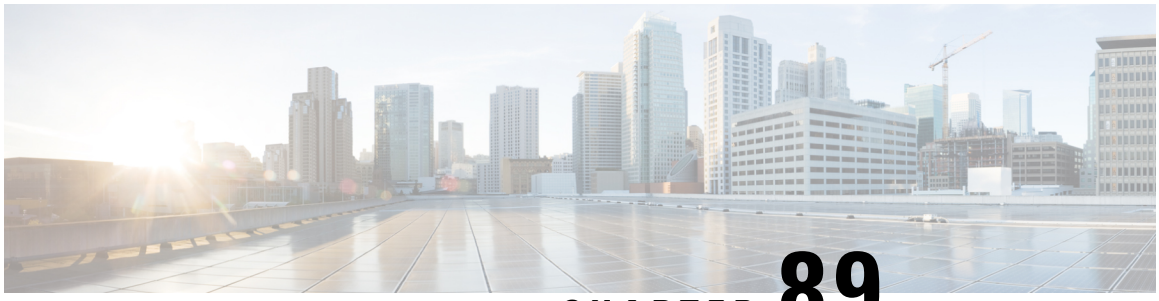
Step 8 (Optional.) Click **Logging** and decide whether you want to see syslog events for matching traffic.

Step 9 Click **Add** to add the rule.

Take the time to move the rule to the right location in the policy if it is not already there.

Click **Save** to save the policy.

The next time you deploy the configuration, the rule will be applied to the target devices.



CHAPTER 89

DNS Policies for Security Intelligence

The following topics explain DNS policies, DNS rules, and how to deploy DNS policies to managed devices.

- DNS policies, on page 2557
- Cisco umbrella DNS policies, on page 2558
- DNS policy components, on page 2559
- License requirements for DNS policies, on page 2560
- Requirements and prerequisites for DNS policies, on page 2560
- Managing DNS and Umbrella DNS Policies, on page 2560
- DNS rules, on page 2562
- DNS rules, on page 2569
- Requirement: deploy DNS policy configuration changes, on page 2575
- Cisco umbrella DNS policies, on page 2575

DNS policies

A DNS policy is a security configuration that

- blocks traffic based on the domain name requested by a client, using a Security Intelligence Block list
- uses Cisco-provided domain name intelligence to filter traffic
- supports custom lists and feeds of domain names tailored to your deployment, and
- requires association with an access control policy for deployment to managed devices.

DNS policy behavior

Traffic on a DNS policy Block list is immediately blocked and therefore is not subject to any further inspection—not for intrusions, exploits, malware, and so on, but also not for network discovery. You can use a Security Intelligence Do Not Block list to override a Block list and force access control rule evaluation, and, recommended in passive deployments, you can use a "monitor-only" setting for Security Intelligence filtering. This allows the system to analyze connections that would have been blocked by a Block list, but also logs the match to the Block list and generates an end-of-connection Security Intelligence event.

Note DNS-based Security Intelligence may not work as intended for a domain name unless the DNS server deletes a domain cache entry due to expiration, or a client's DNS cache or the local DNS server's cache is cleared or expires.

You configure DNS-based Security Intelligence using a DNS policy and associated DNS rules. To deploy it to your devices, you must associate your DNS policy with an access control policy, then deploy your configuration to managed devices.

Cisco umbrella DNS policies

A Cisco Umbrella DNS policy is a security mechanism that

- redirects DNS queries to Cisco Umbrella for domain validation and policy enforcement,
- allows or blocks requests based on domain names and applies DNS-based security policies, and
- provides two lines of protection through local DNS inspection and cloud-based Cisco Umbrella policies.

DNS query redirection process

Cisco Umbrella DNS Connection in the management center helps to redirect DNS queries to Cisco Umbrella. This allows Cisco Umbrella to validate requests, allow or block them based on the domain names, and apply DNS-based security policy on the request. If you use Cisco Umbrella, you must configure the Cisco Umbrella Connection (**Integration** > **Other Integrations** > **Cloud Services** and choose the **Cisco Umbrella Connection**) to redirect DNS queries to Cisco Umbrella.

The Umbrella Connector is part of the system's DNS inspection. If your existing DNS inspection policy map decides to block or drop a request based on your DNS inspection settings, the request is not forwarded to Cisco Umbrella. Thus, you have two lines of protection:

- Your local DNS inspection policy
- Your Cisco Umbrella cloud-based policy

When redirecting DNS lookup requests to Cisco Umbrella, the Umbrella Connector adds an EDNS (Extension mechanisms for DNS) record. An EDNS record includes the device identifier information, organization ID, and client IP address. Your cloud-based policy can use those criteria to control access in addition to the reputation of the FQDN. You can also elect to encrypt the DNS request using DNSCrypt to ensure the privacy of usernames and internal IP addresses.

To redirect DNS requests from the management center to Cisco Umbrella:

1. Configure the Cisco Umbrella connection settings.
2. Create and configure an Umbrella DNS policy.
3. Associate the Umbrella DNS policy with an access control policy.
4. Deploy the changes.

For detailed information about how to set up the Umbrella DNS Connector in the management center, see Configuring the Umbrella DNS Connector for Cisco Secure Firewall Management Center.

DNS policy components

DNS policy components are configuration elements that

- enable blocking of connections based on domain name using a Block list
- exempt connections from blocking using a Do Not Block list, and
- provide granular network traffic handling through configurable rules and settings.

DNS policy configuration options

These configurations can be changed after creating a DNS policy:

- **Name and Description**: Each DNS policy must have a unique name. A description is optional.
- **Rules**: Rules provide a granular method of handling network traffic based on the domain name. Rules in a DNS policy are numbered, starting at 1. The system matches traffic to DNS rules in top-down order by ascending rule number.

When you create a DNS policy, the system populates it with a default Global Do-Not-Block List for DNS rule and a default Global Block List for DNS rule. Both rules are fixed to the first position in their respective categories. You cannot modify these rules, but you can disable them.



Note If multitenancy is enabled for your Cloud-Delivered Firewall Management Center, the system is organized into a hierarchy of domains, including ancestor and descendant domains. These domains are distinct and separate from the domain names used in DNS management.

A descendant list contains the domains on the Block or Do Not Block lists of system subdomain users. From an ancestor domain, you cannot view the contents of descendant lists. If you do not want subdomain users to add domains to a Block or Do Not Block list:

- disable the descendant list rules, and
- enforce Security Intelligence using the access control policy inheritance settings

The system evaluates rules in this order:

1. Global Do-Not-Block List for DNS rule (if enabled)
2. Descendant DNS Do-Not-Block Lists rule (if enabled)
3. Rules with a Do Not Block action
4. Global Block List for DNS rule (if enabled)
5. Descendant DNS Block Lists rule (if enabled)
6. Rules with an action other than Do Not Block

Usually, the system handles DN-based network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic. If no DNS rules match the traffic, the system continues evaluating the traffic based on the associated access control policy's rules. DNS rule conditions can be simple or complex.

License requirements for DNS policies

This reference provides the license requirements necessary for implementing DNS policies across different Firepower system configurations.

Threat Defense License

IPS

Classic License

Protection

Requirements and prerequisites for DNS policies

This reference provides the requirements and prerequisites needed to configure and implement DNS policies, including model support, domain compatibility, and required user roles.

Model support

Any

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Important You must apply the Network Discovery policy on the device for a successful DNS validation on the traffic.

Managing DNS and Umbrella DNS Policies

Use the DNS Policy page to create, edit, delete, copy, and compare custom DNS and Umbrella DNS policies for comprehensive network security management.

Use the DNS Policy page (**Policies** > **Access Control heading** > **DNS**) to manage custom DNS and Umbrella DNS policies.

In addition to custom policies that you create, the system provides the Default DNS Policy and Default Umbrella DNS Policy. The default DNS policy uses the default Block list and Do Not Block list. You can edit and use this system-provided custom policies.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 Manage your DNS policy:

- Compare—To compare DNS policies, click **Compare Policies** and proceed as described in Compare policies, on page 219.
- Copy—To copy a DNS policy, click **Copy** ( ) and proceed as described in Edit DNS policies, on page 2562.
- Create—To create a new Umbrella DNS policy, click **New Policy** > **Umbrella DNS Policy** and proceed as described in Create an umbrella DNS policy, on page 2579.
- Delete—To delete a DNS or Umbrella DNS policy, click **Delete** ( ), then confirm you want to delete the policy.
- Edit—To modify an existing DNS policy, click **Edit** ( )) and proceed as described in Edit DNS policies, on page 2562. To modify an existing Umbrella DNS policy, click **Edit** ( )) and proceed as described in Edit Umbrella DNS policies and rules, on page 2579.

You have successfully managed your DNS policies. The policies are now available for deployment to your network devices for DNS-based security filtering.

Create basic DNS policies

Creating a basic DNS policy establishes a foundation for DNS traffic filtering and management that you can customize to meet your organization's security requirements.

When you create a new DNS policy, it contains default settings. You must then edit it to customize the behavior.

Follow these steps to create a basic DNS policy:

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 Click **Add DNS Policy** > **DNS Policy**.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

Step 4 Click **Save**.

The DNS policy is created with default settings and appears in the DNS policies list.

What to do next

Configure the policy. See Edit DNS policies, on page 2562.

Edit DNS policies

Edit DNS policies to update policy configurations, modify rule settings, and maintain network security requirements.

Only one person should edit a DNS policy at a time, using a single browser window. If multiple users attempt to save the same policy, only the first set of saved changes are retained.

To protect the privacy of your session, after thirty minutes of inactivity on the policy editor, a warning appears. After sixty minutes, the system discards your changes.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 Click **Edit** ( ) next to the DNS policy you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your DNS policy:

- Name and Description—To change the name or description, click the field and type the new information.
- Rules—To add, categorize, enable, disable, or otherwise manage DNS rules, click **Rules** and proceed as described in Create and edit DNS rules, on page 2563.

Step 4 Click **Save**.

The DNS policy is updated with your changes and ready for deployment.

What to do next

- Optionally, further configure the new policy as described in *Logging Connections with Security Intelligence* in the *Cisco Secure Firewall Management Center Administration Guide*.
- Deploy configuration changes; see Deploy configuration changes, on page 202.

DNS rules

A DNS rule is a Security Intelligence component that

- handles traffic based on the domain name requested by a host,
- operates after any traffic decryption and before access control evaluation, and
- matches traffic in the order specified, with the system handling network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic.

DNS rule components

In addition to its unique name, each DNS rule has these basic components:

- **State**: By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.
- **Position**: Rules in a DNS policy are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.
- **Conditions**: Conditions specify the specific traffic the rule handles. A DNS rule must contain a DNS feed or list condition, and can also match traffic by security zone, network, dynamic attributes, or VLAN.
- **Action**: A rule's action determines how the system handles matching traffic.

A rule's action determines how the system handles matching traffic:

- Traffic with a **Do Not Block** action is allowed, subject to further access control inspection.
- Monitored traffic is subject to further evaluation by remaining rules on the DNS Block list. If the traffic does not match a DNS Block list rule, it is inspected with access control rules. The system logs a Security Intelligence event for the traffic.
- Traffic on a Block list is dropped without further inspection. You can also return a Domain Not Found response, or redirect the DNS query to a sinkhole server.

Related Topics

Security intelligence, on page 2027

Create and edit DNS rules

Create and edit DNS rules to define how the system handles DNS queries by specifying actions and conditions for domain name resolution.

In a DNS policy, you can add up to a total of 32767 DNS lists to the Block list and Do Not Block list rules; that is, the number of lists in the DNS policy cannot exceed 32767.

Follow these steps to create and edit DNS rules:

Procedure

Step 1 In the DNS policy editor, you have the following options:

- To add a new rule, click **Add DNS Rule**.
- To edit an existing rule, click **Edit** ( ).

Step 2 Enter a **Name**.

Step 3 Configure the rule components, or accept the defaults:

- Action—Choose a rule **Action**; see DNS rule actions, on page 2565.
- Conditions—Configure the rule's conditions; see DNS rule conditions, on page 2566.
- Enabled—Specify whether the rule is **Enabled**.

Step 4 Click **Save**.

The DNS rule is created or updated with the specified configuration and is ready for deployment.

What to do next

- Deploy configuration changes; see Deploy configuration changes, on page 202.

DNS rule management

DNS rule management is a policy administration feature that enables you to add, edit, move, enable, disable, delete, and otherwise manage DNS rules within your policy.

DNS rule management interface

The **Rules** tab of the DNS policy editor provides the interface for DNS rule management operations.

For each rule, the policy editor displays this information:

- Rule name
- Summary of rule conditions
- Rule action

Additional icons represent **Warning** ( ), **Error** ( ), and other important **Information** ( ). Disabled rules are dimmed and marked `(disabled)` beneath the rule name.

Enable and disable DNS rules

Enable or disable DNS rules to control whether the system uses them to evaluate network traffic and generate warnings and errors.

When you create a DNS rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in a DNS policy, disabled rules are dimmed, although you can still modify them. Note that you can also enable or disable a DNS rule using the DNS rule editor.

Follow these steps to enable or disable DNS rules:

Procedure

Step 1 In the DNS policy editor, right-click the rule and choose a rule state.

Step 2 Click **Save**.

The DNS rule state is updated. Disabled rules appear dimmed in the policy rule list but can still be modified.

What to do next

- Deploy configuration changes.

DNS rule order evaluation

DNS rule order evaluation is a traffic processing mechanism that

- matches traffic to DNS rules in top-down order by ascending rule number,
- handles network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic, and
- applies different continuation behavior based on rule type.

Rule evaluation behavior

Rules in a DNS policy are numbered, starting at 1. The system behavior differs based on rule type:

- For Monitor rules, the system logs the traffic, then continues evaluating traffic against lower-priority DNS Block list rules.
- For non-Monitor rules, the system does **not** continue to evaluate traffic against additional, lower-priority DNS rules after that traffic matches a rule.

Rule order precedence follows this hierarchy:

- The Global Do-Not-Block List for DNS is always first, and takes precedence over all other rules.
- The Do-Not-Block List section precedes the Block List section; Do-Not-Block List rules always take precedence over other rules.
- The Global Block List for DNS is always first in the Block List section, and takes precedence over all other Monitor and Block list rules.
- The Block List section contains Monitor and Block list rules.
- When you first create a DNS rule, the system positions it last in the Do-Not-Block List section if you assign a **Do Not Block** action, or last in the Block List section if you assign any other action.

You can drag and drop rules to reorder them.

DNS rule actions

A DNS rule action is a configuration setting that

- determines handling for matching traffic, governing whether the system will block, not block, or monitor traffic based on rule conditions,
- controls logging behavior, determining when and how details about matching traffic are recorded, and
- impacts action prioritization when configured with TID.

DNS rule action types

DNS rules support these action types:

- **Do Not Block**: Allows traffic to pass to the next phase of inspection, which is access control rules. The system does not log *Do Not Block list* matches. Logging of these connections depends on their eventual disposition.
- **Monitor**: Forces connection logging; matching traffic is neither immediately allowed nor blocked. Traffic is matched against additional rules to determine whether to permit or deny it. The first non-Monitor DNS rule matched determines whether the system blocks the traffic. If there are no additional matching rules, the traffic is subject to access control evaluation. For connections monitored by a DNS policy, the system logs end-of-connection Security Intelligence and connection events to the Cloud-Delivered Firewall Management Center database.
- **Drop**: Drops the traffic without further inspection.
- **Domain Not Found**: Returns a non-existent internet domain response to the DNS query, which prevents the client from resolving the DNS request.
- **Sinkhole**: Returns a sinkhole object's IPv4 or IPv6 address in response to the DNS query (A and AAAA records only). The sinkhole server can log, or log and block, follow-on connections to the IP address. If you configure a **Sinkhole** action, you must also configure a sinkhole object.

Logging behavior varies based on the action type:

- For a connection blocked based on the **Drop** or **Domain Not Found** actions, the system logs beginning-of-connection Security Intelligence and connection events. Because blocked traffic is immediately denied without further inspection, there is no unique end of connection to log.
- For a connection blocked based on the **Sinkhole** action, logging depends on the sinkhole object configuration. If you configure your sinkhole object to only log sinkhole connections, the system logs end-of-connection connection events for the follow-on connection. If you configure your sinkhole object to log and block sinkhole connections, the system logs beginning-of-connection connection events for the follow-on connection, then blocks that connection.

DNS rule conditions

A DNS rule condition is a traffic identification criterion that

- identifies the type of traffic that rule handles,
- can be simple or complex, and
- must include a DNS feed or list condition within a DNS rule.

DNS rule condition requirements

You can also optionally control traffic by security zone, network, dynamic attributes, or VLAN.

When adding conditions to a DNS rule:

- If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion.

- You can configure multiple conditions per rule. Traffic must match **all** the conditions in the rule for the rule to apply to traffic. For example, a rule with a DNS feed or list condition and network condition but no VLAN tag condition evaluates traffic based on the domain name and source or destination, regardless of any VLAN tagging in the session.
- For each condition in a rule, you can add up to 50 criteria. Traffic that matches **any** of a condition's criteria satisfies the condition. For example, you can use a single rule to block traffic based on up to 50 DNS lists and feeds.

Security zone rule conditions

A security zone rule condition is a network traffic control mechanism that

- segments your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices,
- controls or decrypts traffic by its source and destination security zones, and
- requires all interfaces in a zone to be of the same type (all inline, passive, switched, or routed).

Zone condition traffic matching

If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

All zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Network rule conditions

A network condition is a traffic matching criterion that

- controls or decrypts traffic based on source and destination IP addresses using inner packet headers,
- enables administrators to specify exact IP addresses or blocks for granular policy enforcement, and
- can be built using predefined network objects or by manually entering IP addresses or address blocks.

Network rule condition restrictions

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

VLAN tag rule conditions

VLAN tags rule conditions are filtering mechanisms that

- control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic
- use the innermost VLAN tag to filter VLAN traffic, with the exception of prefilter policy which uses the outermost VLAN tag, and
- apply only to inline sets and do not match traffic on firewall interfaces.

VLAN tag rule conditions configuration details

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note this Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

DNS policy rule conditions

A DNS policy rule condition is a DNS rule component that

- allows you to control traffic if a DNS list, feed, or category contains the domain name requested by the client,
- must be defined in a DNS rule, and
- applies the configured rule action to traffic regardless of whether you add a global or custom Block or Do Not Block list.

DNS condition behavior

The system applies the configured rule action to the traffic regardless of the list type. For example, if you add the Global Do Not Block List to a rule, and configure a **Drop** action, the system blocks all traffic that should have been allowed to pass to the next phase of inspection.

Dynamic attributes rule conditions

A dynamic attributes rule condition is a DNS rule matching mechanism that

- uses dynamic objects containing IP addresses or endpoint device type objects for source or destination matching,
- uses Security Group Tag (SGT) objects containing manually defined tags or ISE-defined tags for source-only matching, and
- applies logical operations where objects of the same type are ORed together and objects of different types are ANDed together.

Dynamic attributes configuration and behavior

You can use these types of dynamic attributes to match connections in DNS rules:

- (Source or destination.) Dynamic objects, which contain IP addresses. Endpoint device type objects are source only. For more information, see Dynamic objects, on page 2966 and the chapter on Dynamic Attributes Connector.
- (Source only.) Security Group Tag (SGT) objects, which contain tags either manually defined or defined through ISE. For more information, see Source and destination security group tag (SGT) matching, on page 2374 and Security group tags, on page 2972.

When you configure dynamic attributes for a DNS rule, objects of the same type in the same source or destination list are ORed together and objects of different types are ANDed together. For example, if you select both a security group tag, and a dynamic object that lists IP addresses, the rule matches if traffic with the tag originates from (or is destined to) one of those IP addresses.

Initially, all **Dynamic Objects** and **Security Group Tags** are listed when you open the **Dynamic Attributes** tab. You can deselect an option to remove those objects from the list. You can also start typing in the search box to find the object you want.

DNS rules

A DNS rule is a network security mechanism that controls and filters DNS queries and responses within network traffic

Control traffic based on DNS and security zone

Control traffic flow by implementing DNS rules with security zone conditions to manage network access and security policies.

Zone conditions in DNS rules allow you to control traffic by its source security zone. A *security zone* is a grouping of one or more interfaces, which may be located across multiple devices.

Procedure

Step 1 In the DNS rule editor, click **Zones**.

Step 2 Find and select the zones you want to add from the **Available Zones**. To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Step 3 Click to select a zone, or right-click and then select **Select All**.

Step 4 Click **Add to Source**, or drag and drop.

Step 5 Click the **DNS** tab and add the lists or feeds that include the DNS names you are controlling. For more information, see Control traffic based on DNS list or feed, on page 2572.

Step 6 Save or continue editing the rule.

The DNS rule is configured with security zone conditions to control traffic based on source security zones and DNS names.

What to do next

- Deploy configuration changes.

Control traffic based on DNS and network

Control traffic by explicitly specifying the source IP addresses for the traffic you want to manage through DNS rules.

Network conditions in DNS rules allow you to control traffic by its source IP address. You can explicitly specify the source IP addresses for the traffic you want to control.

Procedure

Step 1 In the DNS rule editor, click **Networks**.

Step 2 Find and select the networks you want to add from the **Available Networks**, as follows:

- To add a network object on the fly, which you can then add to the condition, click **Add** (+) above the **Available Networks** list and proceed as described in Create a network object, on page 2985.
- To search for network objects to add, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

Step 3 Click **Add to Source**, or drag and drop.

Step 4 Add any source IP addresses or address blocks that you want to specify manually. Click the **Enter an IP address** prompt below the **Source Networks** list; then type an IP address or address block and click **Add**.

Step 5 Click the **DNS** tab and add the lists or feeds that include the DNS names you are controlling. For more information, see Control traffic based on DNS list or feed, on page 2572.

Step 6 Save or continue editing the rule.

The DNS rule is configured with network conditions to control traffic based on the specified source IP addresses and DNS names.

What to do next

- Deploy configuration changes.

Control traffic based on DNS and VLAN

This task allows you to control VLAN-tagged traffic by configuring VLAN conditions in DNS rules, enabling precise traffic management based on DNS names and VLAN identifiers.

VLAN conditions in DNS rules allow you to control VLAN-tagged traffic. The system uses the innermost VLAN tag to identify a packet by VLAN.

When you build a VLAN-based DNS rule condition, you can manually specify VLAN tags. Alternately, you can configure VLAN conditions with VLAN tag *objects*, which are reusable and associate a name with one or more VLAN tags.

Procedure

Step 1 In the DNS rule editor, select **VLAN Tags**.

Step 2 Find and select the VLANs you want to add from the **Available VLAN Tags**, as follows:

- To add a VLAN tag object on the fly, which you can then add to the condition, click **Add** (+) above the Available VLAN Tags list and proceed as described in Create a VLAN tag object, on page 3057.
- To search for VLAN tag objects and groups to add, click the **Search by name or value** prompt above the **Available VLAN Tags** list, then type either the name of the object, or the value of a VLAN tag in the object. The list updates as you type to display matching objects.

Step 3 Click **Add to Rule**, or drag and drop.

Step 4 Add any VLAN tags that you want to specify manually. Click the **Enter a VLAN Tag** prompt below the **Selected VLAN Tags** list; then type a VLAN tag or range and click **Add**. You can specify any VLAN tag from 1 to 4094; use a hyphen to specify a range of VLAN tags.

Step 5 Click the **DNS** tab and add the lists or feeds that include the DNS names you are controlling. For more information, see Control traffic based on DNS list or feed, on page 2572.

Step 6 Save or continue editing the rule.

The DNS rule is configured with VLAN conditions, allowing the system to control traffic based on both DNS names and VLAN tags.

What to do next

- Deploy configuration changes.

Control traffic based on DNS list or feed

Control network traffic by applying DNS-based filtering rules using predefined or custom DNS lists and threat intelligence feeds.

Use DNS lists and feeds in security policies to block or allow traffic based on domain reputation and threat intelligence data.

Procedure

Step 1 In the DNS rule editor, click **DNS**.

Step 2 Find and select the DNS lists and feeds you want to add from the **DNS Lists and Feeds**, as follows:

- To add a DNS list or feed on the fly, which you can then add to the condition, click **Add** (+) above the **DNS Lists and Feeds** list and proceed as described in Create custom security intelligence feeds, on page 3030.
- To search for DNS lists, feeds, or categories to add, click the **Search by name or value** prompt above the **DNS Lists and Feeds** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.
- For descriptions of the system-provided threat categories, see Security intelligence categories, on page 2034.

Step 3 Click **Add to Rule**, or drag and drop.

Step 4 Save or continue editing the rule.

The DNS lists and feeds are now configured in your rule and will filter traffic based on the selected criteria.

What to do next

- Deploy configuration changes.

Controlling traffic based on security group tag or dynamic attributes

This task allows you to control traffic using dynamic attributes conditions in DNS rules based on security group tags or other dynamic attributes.

Dynamic Attributes conditions in DNS rules allow you to control traffic based on security group tag or other dynamic attributes. To use these attributes, you must apply security group tags to traffic in your network, or create the dynamic attribute objects you need using the API or the Cisco Secure Dynamic Attributes Connector. For more information about enabling dynamic attributes, see Dynamic attributes rule conditions, on page 2569.

When you configure dynamic attributes for a DNS rule, objects of the same type in the same source or destination list are ORed together and objects of different types are ANDed together. For example, if you select both a security group tag, and a dynamic object that lists IP addresses, the rule matches if traffic with the tag originates from (or is destined to) one of those IP addresses.

Procedure

Step 1 In the DNS rule editor, click **Dynamic Attributes**.

Step 2 Select the objects you want to use, and click either **Add Sources** or **Add Destinations**.

Initially, all **Dynamic Objects** and **Security Group Tags** are listed when you open the **Dynamic Attributes** tab. You can deselect an option to remove those objects from the list. You can also start typing in the search box to find the object you want.

You can use these objects to identify the source of the connection, the destination, or both.

Note
You can use SGT and endpoint device type objects as source criteria only.

Step 3 Click the **DNS** tab and add the lists or feeds that include the DNS names you are controlling. For more information, see Control traffic based on DNS list or feed, on page 2572.

Step 4 Save or continue editing the rule.

You have successfully configured dynamic attributes for DNS rules to control traffic based on security group tags or other dynamic attributes.

What to do next

- Deploy configuration changes.

Use a DNS sinkhole to enforce content restriction

This task configures a DNS sinkhole to redirect traffic to the Google SafeSearch Virtual IP Address (VIP), which imposes content filters on Google and YouTube search results.

Typically, a DNS sinkhole directs traffic away from a particular target. This procedure describes how to configure a DNS sinkhole to redirect traffic to the Google SafeSearch Virtual IP Address (VIP), which imposes content filters on Google and YouTube search results.

Because Google SafeSearch uses a single IPv4 address for the VIP, hosts must use IPv4 addressing.



Caution If your network includes proxy servers, this content restriction method is not effective unless you position your Firewall Threat Defense devices between the proxy servers and the Internet.

This procedure enforces content restriction for Google searches only. It does not apply to any other search engine.

Before you begin

This procedure applies to Firewall Threat Defense only, and requires the IPS license.

Follow these steps to use a DNS sinkhole to enforce content restriction:

Procedure

Step 1 Obtain a list of supported Google domains via the following URL: https://www.google.com/supported_domains.

Step 2 Create a custom DNS list on your local computer, and add the following entries:

- To enforce Google SafeSearch, add an entry for each supported Google domain.
- To enforce YouTube Restricted Mode, add a "youtube.com" entry.

The custom DNS list must be in text file (.txt) format. Each line of the text file must specify an individual domain name, stripped of any leading periods. For example, the supported domain ".google.com" must appear as "google.com".

Step 3 Upload the custom DNS list to the Cloud-Delivered Firewall Management Center; see Upload new custom security intelligence lists to the Cloud-Delivered Firewall Management Center, on page 3032.

Step 4 Determine the IPv4 address for the Google SafeSearch VIP. For example, run **nslookup** on forcesafesearch.google.com.

Step 5 Create a sinkhole object for the SafeSearch VIP; see Create a sinkhole object, on page 3033.

Use the following values for this object:

- IPv4 Address—Enter the SafeSearch VIP address.
- IPv6 Address—Enter the IPv6 loopback address (**::1**).
- Log Connections to Sinkhole—Click Log Connections.
- Type—Choose **None**.

Step 6 Create a basic DNS policy; see Create basic DNS policies, on page 2561.

Step 7 Add a DNS rule for the sinkhole; see Create and edit DNS rules, on page 2563.

For this rule:

- Check the **Enabled** check box.
- Choose **Sinkhole** from the **Action** drop-down list.
- Choose the sinkhole object you created from the **Sinkhole** drop-down list.
- Add the custom DNS list you created to the **Selected Items** list on **DNS**.
- (Optional) Choose a network in **Networks** to limit content restriction to specific users. For example, if you want to limit content restriction to student users, assign students to a different subnet than faculty, and specify that subnet in this rule.

Step 8 Associate the DNS policy with an access control policy; see Associate prefilter, decryption, and identity policies with an access control, on page 2016.

Step 9 Deploy configuration changes.

The DNS sinkhole is configured to enforce content restriction by redirecting traffic to the Google SafeSearch VIP, which filters Google and YouTube search results.

Requirement: deploy DNS policy configuration changes

After you finish updating your DNS policy configuration, you must deploy it as part of access control configuration.

- Associate your DNS policy with an access control policy, as described in Configure security intelligence, on page 2030.
- Deploy configuration changes.

Cisco umbrella DNS policies

A Cisco Umbrella DNS policy is a security mechanism that

- redirects DNS queries to Cisco Umbrella for domain validation and policy enforcement,
- allows or blocks requests based on domain names and applies DNS-based security policies, and
- provides two lines of protection through local DNS inspection and cloud-based Cisco Umbrella policies.

DNS query redirection process

Cisco Umbrella DNS Connection in the management center helps to redirect DNS queries to Cisco Umbrella. This allows Cisco Umbrella to validate requests, allow or block them based on the domain names, and apply DNS-based security policy on the request. If you use Cisco Umbrella, you must configure the Cisco Umbrella Connection (**Integration** > **Other Integrations** > **Cloud Services** and choose the **Cisco Umbrella Connection**) to redirect DNS queries to Cisco Umbrella.

The Umbrella Connector is part of the system's DNS inspection. If your existing DNS inspection policy map decides to block or drop a request based on your DNS inspection settings, the request is not forwarded to Cisco Umbrella. Thus, you have two lines of protection:

- Your local DNS inspection policy
- Your Cisco Umbrella cloud-based policy

When redirecting DNS lookup requests to Cisco Umbrella, the Umbrella Connector adds an EDNS (Extension mechanisms for DNS) record. An EDNS record includes the device identifier information, organization ID, and client IP address. Your cloud-based policy can use those criteria to control access in addition to the reputation of the FQDN. You can also elect to encrypt the DNS request using DNSCrypt to ensure the privacy of usernames and internal IP addresses.

To redirect DNS requests from the management center to Cisco Umbrella:

1. Configure the Cisco Umbrella connection settings.
2. Create and configure an Umbrella DNS policy.
3. Associate the Umbrella DNS policy with an access control policy.
4. Deploy the changes.

For detailed information about how to set up the Umbrella DNS Connector in the management center, see Configuring the Umbrella DNS Connector for Cisco Secure Firewall Management Center.

Redirect DNS requests to Cisco umbrella

Redirect DNS requests from the device to Cisco Umbrella for enhanced security and DNS filtering capabilities.

Use this procedure when you need to configure DNS redirection to Cisco Umbrella using the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Ensure that you meet the prerequisites.

Prerequisites for configure the umbrella DNS connector, on page 2576

Step 2 Configure the Cisco Umbrella connection settings.

Configure Cisco umbrella connection settings, on page 2578

Step 3 Create an Umbrella DNS policy.

Create an umbrella DNS policy, on page 2579

Step 4 Configure the Umbrella DNS policy.

Edit Umbrella DNS policies and rules, on page 2579

Step 5 Associate the Umbrella DNS policy with an access control policy.

Associate the umbrella DNS policy with an access control policy, on page 2580

DNS requests from the device are now redirected to Cisco Umbrella for processing and security filtering.

Prerequisites for configure the umbrella DNS connector

This reference provides the prerequisites that must be met before configuring the Umbrella DNS Connector, including minimum platform requirements, account setup, certificate management, and network connectivity requirements.

Table 137: Minimum supported platforms

Product	Version
Secure Firewall Threat Defense	6.6 and above
Cloud-Delivered Firewall Management Center	7.2 and above

- Establish an account with Cisco Umbrella at https://umbrella.cisco.com, and log into Umbrella at http://login.umbrella.com.

- Import the CA certificate from the Cisco Umbrella server to the Cloud-Delivered Firewall Management Center. In Cisco Umbrella, choose **Deployments** > **Configuration** > **Root Certificate** and download the certificate.

  You must import the root certificate to establish the HTTPS connection with the Cisco Umbrella registration server. The certificate needs to be trusted for SSL server validation, which is a non-default option in the Cloud-Delivered Firewall Management Center. Copy and paste this certificate for the device in the Cloud-Delivered Firewall Management Center(**Devices** > **Certificates**).

  ```
  MIIE6jCCA9KgAwIBAgIQCjUI1VwpKwF9+K1lwA/35DANBgkqhkiG9w0BAQsFADBhMQswCQYDVQQG
  EwJVUzEVMBMGA1UEChMMRGlnaUNlcnQgSW5jMRkwFwYDVQQLExB3d3cuZGlnaWNlcnQuY29tMSAw
  HgYDVQQDExdEaWdpQ2VydCBHbG9iYWwgUm9vdCBDQTAeFw0yMDA5MjQwMDAwMDBaFw0zMDA5MjMy
  MzU5NTlaME8xCzAJBgNVBAYTAlVTMRUwEwYDVQQKEwxEaWdpQ2VydCBJbmMxKTAnBgNVBAMTIERp
  Z2lDZXJ0IFRMUyBSU0EgU0hBMjU2IDIwMjAgQ0ExMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIB
  CgKCAQEAwUuzZUdwvN1PWNvsnO3DZuUfMRNUrUpmRh8sCuxkB+Uu3Ny5CiDt3+PE0J6aqXodgojl
  EVbbHp9YwlHnLDQNLtKS4VbL8Xlfs7uHyiUDe5pSQWYQYE9XE0nw6Ddng9/n00tnTCJRpt8OmRDt
  V1F0JuJ9x8piLhMbfyOIJVNvwTRYAIuE//i+p1hJInuWraKImxW8oHzf6VGo1bDtN+I2tIJLYrVJ
  muzHZ9bjPvXj1hJeRPG/cUJ9WIQDgLGBAfr5yjK7tI4nhyfFK3TUqNaX3sNk+crOU6JWvHgXjkkD
  Ka77SU+kFbnO8lwZV21reacroicgE7XQPUDTITAHk+qZ9QIDAQABo4IBrjCCAaowHQYDVR0OBBYE
  FLdrouqoqoSMeeq02g+YssWVdrn0MB8GA1UdIwQYMBaAFAPeUDVW0Uy7ZvCj4hsbw5eyPdFVMA4G
  A1UdDwEB/wQEAwIBhjAdBgNVHSUEFjAUBggrBgEFBQcDAQYIKwYBBQUHAwIwEgYDVR0TAQH/BAgw
  BgEB/wIBADB2BggrBgEFBQcBAQRqMGgwJAYIKwYBBQUHMAGGGGh0dHA6Ly9vY3NwLmRpZ2ljZXJ0
  LmNvbTBABggrBgEFBQcwAoY0aHR0cDovL2NhY2VydHMuZGlnaWNlcnQuY29tL0RpZ2lDZXJ0R2xv
  YmFsUm9vdENBLmNydDB7BgNVHR8EdDByMDegNaAzhjFodHRwOi8vY3JsMy5kaWdpY2VydC5jb20v
  RGlnaUNlcnRHbG9iYWxSb290Q0EuY3JsMDegNaAzhjFodHRwOi8vY3JsNC5kaWdpY2VydC5jb20v
  RGlnaUNlcnRHbG9iYWxSb290Q0EuY3JsMDAGA1UdIAQpMCcwBwYFZ4EMAQEwCAYGZ4EMAQIBMAgG
  BmeBDAECAjAIBgZngQwBAgMwDQYJKoZIhvcNAQELBQADggEBAHert3onPa679n/gWlbJhKrKW3EX
  3SJH/E6f7tDBpATho+vFScH90cnfjK+URSxGKqNjOSD5nkoklEHIqdninFQFBstcHL4AGw+oWv8Z
  u2XHFq8hVt1hBcnpj5h232sb0HIMULkwKXq/YFkQZhM6LawVEWwtIwwCPgU7/uWhnOKK24fXSuhe
  50gG66sSmvKvhMNbg0qZgYOrAKHKCjxMoiWJKiKnpPMzTFuMLhoClw+dj20tlQj7T9rxkTgl4Zxu
  YRiHas6xuwAwapu3r9rxxZf+ingkquqTgLozZXq8oXfpf2kUCwA/d5KxTVtzhwoT0JzI8ks5T1KE
  SaZMkE4f97Q=
  ```

  When you add the certificate in the management center, ensure that you check the **CA Only** check box.

- Install the certificate on the device.
- Obtain this data from Umbrella:
  - Organization ID
  - Network Device Key
  - Network Device Secret
  - Legacy Network Device Token
- Ensure that the Cloud-Delivered Firewall Management Center is connected to the internet.
- Ensure that the base license with the export-controlled feature option is enabled in the Cloud-Delivered Firewall Management Center.
- Ensure that the DNS server is configured to resolve api.opendns.com.
- Ensure that the Cloud-Delivered Firewall Management Center can resolve management.api.umbrella.com for policy configuration.
- Configure the Firewall Threat Defense route to api.opendns.com.

Configure Cisco umbrella connection settings

Configure the token and settings required to register the device with Cisco Umbrella and enable DNS security services.

The Cisco Umbrella Connection settings define the token that is needed to register the device with Cisco Umbrella.

Before you begin

Establish an account with Cisco Umbrella https://umbrella.cisco.com, and then log into Umbrella at https://dashboard.umbrella.com and obtain the required information to establish connection to Cisco Umbrella.

Follow these steps to configure Cisco Umbrella connection settings:

Procedure

Step 1 Choose **Integration** > **Other Integrations** > **Cloud Services** and choose the **Cisco Umbrella Connection**.

Step 2 Obtain the following details and add them to the **General** settings:

- **Organization ID**—A unique number that identifies your organization on Cisco Umbrella. Every Umbrella organization is a separate instance of Umbrella and has its own dashboard. Organizations are identified by their name and their organization ID (Org ID).
- **Network Device Key**—The key to fetch umbrella policy from Cisco Umbrella.
- **Network Device Secret**—The secret to fetch umbrella policy from Cisco Umbrella.
- **Legacy Network Device Token**—An Umbrella Legacy Network Devices API token is issued through the Cisco Umbrella dashboard. Umbrella requires the API token to register a network device.

Step 3 Under **Advanced**, configure the following optional settings:

- **DNSCrypt Public Key**—DNSCrypt authenticates and encrypts the DNS queries between the endpoint and the DNS server. To enable DNScrypt, you can configure the DNScrypt public key for certificate verification. The key is a 32-byte hexadecimal value and is preconfigured to B735:1140:206F:225d:3E2B:d822:D7FD:691e:A1C3:3cc8:D666:8d0c:BE04:bfab:CA43:FB79, which is the public key of the Umbrella Anycast servers.
- **Management Key**—A key to fetch datacenter details from Umbrella cloud for VPN policy.
- **Management Secret**—A secret used to fetch datacenters from Umbrella cloud for VPN.

Step 4 Click **Test Connection**—Test if the Cisco Umbrella Cloud is reachable from the Cloud-Delivered Firewall Management Center. When you provide the required organization ID and network device details, the umbrella connection is created.

Step 5 Click **Save**.

The Cisco Umbrella connection is configured and the device is ready to register with Cisco Umbrella using the specified credentials and settings.

Create an umbrella DNS policy

Create an Umbrella DNS policy to implement DNS-based security controls and filtering for network protection.

Umbrella DNS policies provide DNS security by blocking malicious domains and controlling access to web content through DNS filtering.

Before you begin

Follow these steps to create an Umbrella DNS policy:

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 Click **Add DNS Policy** > **Umbrella DNS Policy**.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

Step 4 Click **Save**.

The Umbrella DNS policy is created and ready for configuration with DNS security rules.

What to do next

Configure the policy. See Edit Umbrella DNS policies and rules, on page 2579.

Edit Umbrella DNS policies and rules

Edit Umbrella DNS policies and rules to modify protection policies, configure bypass domains, enable DNSCrypt encryption, and adjust connection timeout settings to meet your organization's specific DNS security requirements.

Use this procedure when you need to modify existing Umbrella DNS policies and their associated rules. You can update protection policy settings, configure local domain bypasses, enable encryption options, and adjust timeout parameters to optimize DNS security for your network environment.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 In the DNS Policy page, select the Umbrella DNS policy that you want to edit and click **Edit** ( ).

Refresh the Umbrella Protection Policy

If you want to get the latest Umbrella Protection Policy from Cisco Umbrella, click the **Refresh** icon next to **Umbrella Protection Policy Last Updated**.

To configure or modify Umbrella Connection settings for the Management Center, go to **Integration** > **Other Integrations** > **Cloud Services** Choose the **Cisco Umbrella Connection**.

Step 3 In the Umbrella DNS policy editor, select the Umbrella DNS rule and click **Edit** ( ).

Step 4 Configure the rule components, or accept the defaults:

- **Umbrella Protection Policy**—Specify the name of the Cisco Umbrella policy to apply to the device.
- **Bypass Domain**—Specify the name of the local domains for which DNS requests should bypass Cisco Umbrella and instead go directly to the configured DNS servers.

  For example, you can have your internal DNS server resolve all names for the organization's domain name on the assumption that all internal connections are allowed.
- **DNSCrypt**— Enable DNSCrypt to encrypt connections between the device and Cisco Umbrella.

  Enabling DNScrypt starts the key-exchange thread with the Umbrella resolver. The key-exchange thread performs the handshake with the resolver every hour and updates the device with a new secret key. As DNScrypt uses UDP/443, you must ensure that the class map used for DNS inspection includes that port. Note that the default inspection class already includes UDP/443 for DNS inspection.
- **Idle Timeout**—Configure the idle timeout after which a connection from a client to the Umbrella server will be removed if there is no response from the server.

Step 5 Click **Save**.

Your Umbrella DNS policy and rule modifications are saved and ready to be deployed. The updated settings will take effect when the policy is applied to devices.

What to do next

Associate the Umbrella DNS policy with an access control policy. For more information, see Associate the umbrella DNS policy with an access control policy, on page 2580.

Associate the umbrella DNS policy with an access control policy

This task associates an Umbrella DNS policy with an access control policy to enable DNS-based security intelligence filtering before deploying the policy on the device.

Before you deploy the Umbrella DNS policy on the device, you must associate it with an access control policy.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control** and select the access policy to edit.

Step 2 Select **Security Intelligence**.

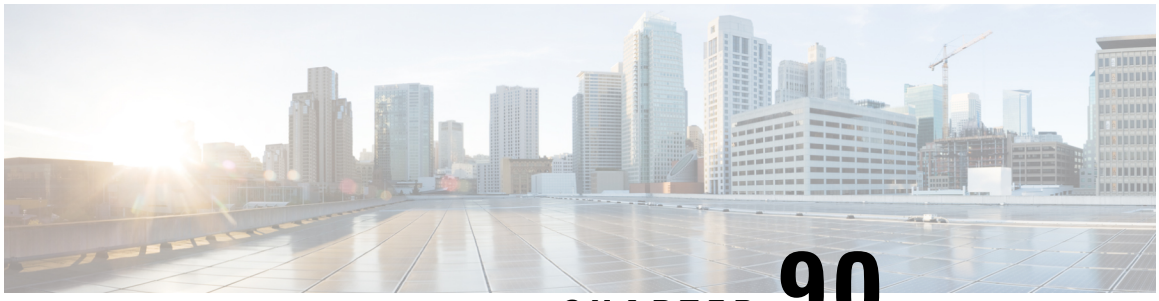
Step 3 From the **Umbrella DNS Policy** drop-down list, select the Umbrella DNS policy.

Step 4 Click **Save**.

The Umbrella DNS policy is now associated with the access control policy and will provide DNS-based security intelligence filtering when deployed.

What to do next

Deploy configuration changes; see Deploy configuration changes, on page 202.



CHAPTER 90

File Policies for Network Malware Protection

The following topics provide an overview of file control, file policies, file rules, Advanced Malware Protection (AMP), cloud connections, and dynamic analysis connections.

- Network malware protection and file policies, on page 2583
- Prerequisites for file policies, on page 2584
- License requirements for file and malware policies, on page 2585
- Best practices for file policies and malware detection, on page 2585
- Configure malware protection, on page 2589
- Cloud connections for malware protection, on page 2594
- File policies and file rules, on page 2597
- Retrospective disposition changes, on page 2613
- File and malware inspection performance and storage options, on page 2613
- Tune file and malware inspection performance and storage, on page 2615
- (Optional) Malware protection with Secure Endpoint, on page 2615

Network malware protection and file policies

Advanced Malware Protection (AMP) for Firepower can detect, capture, track, analyze, log, and optionally block the transmission of malware in network traffic. In the Cloud-Delivered Firewall Management Center web interface, this feature is called malware defense, formerly called AMP for Firepower. Advanced Malware Protection identifies malware using managed devices deployed inline and threat data from the Cisco cloud.

You associate file policies with access control rules that handle network traffic as part of your overall access control configuration.

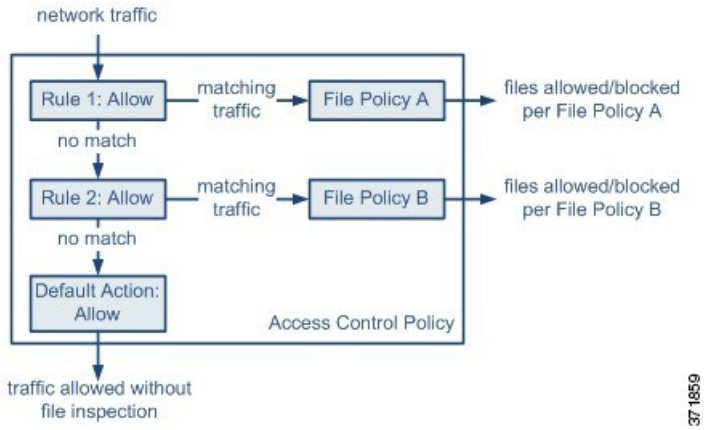
When the system detects malware on your network, it generates file and malware events. To analyze file and malware event data, see the *File/Malware Events and Network File Trajectory* chapter in the *Cisco Secure Firewall Management Center Administration Guide* .

File policies

A file policy is a set of configurations that the system uses to perform malware protection and file control, as part of your overall access control configuration. This association ensures that before the system passes a file in traffic that matches an access control rule's conditions, it first inspects the file. Consider the following diagram of a simple access control policy in an inline deployment.

File policy traffic flow

Consider this diagram of a simple access control policy in an inline deployment.



The policy has two access control rules, both of which use the Allow action and are associated with file policies. The policy's default action is also to allow traffic, but without file policy inspection. In this scenario, traffic is handled as follows:

- Traffic that matches `Rule 1` is inspected by `File Policy A`.
- Traffic that does not match `Rule 1` is evaluated against `Rule 2`. Traffic that matches `Rule 2` is inspected by `File Policy B`.
- Traffic that does not match either rule is allowed; you cannot associate a file policy with the default action.

By associating different file policies with different access control rules, you have granular control over how you identify and block files transmitted on your network.

Prerequisites for file policies

This reference lists the prerequisites necessary for implementing file policies in your network security deployment.

Model support

Any

Secure Firewall 200 supports Malware Analysis through AMP Cloud Lookup only. It does not support the storing of files or the submission of files for analysis. This includes Spero, Local Malware Analysis, File Capture, and Clam AV. Any unsupported feature configured in a policy is automatically turned off when deployed to the device.

Supported domains

Any

User roles

- Admin
- Access Admin

License requirements for file and malware policies

This reference provides the specific license requirements needed to implement various file rule actions and malware detection capabilities in file and malware policies.

To Do This	License Required	File Rule Action
Block or allow all files of a particular type (for example, block all .exe files)	IPS (for Firewall Threat Defense devices) Protection (for Classic devices)	Allow, Block, Block with Reset
Selectively allow or block files based on a judgment that it contains or is likely to contain malware	IPS (for Firewall Threat Defense devices) Protection (for Classic devices) Malware Defense	Malware Cloud Lookup, Block Malware
Store files	IPS (for Firewall Threat Defense devices) Protection (for Classic devices) Malware Defense	Any file rule action with **Store Files** selected

For details about Malware Defense licenses, refer to:

- *Malware Defense Licenses* in the *Cisco Secure Firewall Management Center Administration Guide*

Best practices for file policies and malware detection

In addition to the items described below, follow the steps in Configure malware protection, on page 2589 and referenced topics.

File rule best practices

Consider these guidelines and limitations when configuring file rules:

Prerequisites

- You cannot perform malware analysis on all file types detected by the system. After you select values from the **Application Protocol**, **Direction of Transfer**, and **Action** drop-down lists, the system constrains the list of file types.
- If a file rule is configured with a **Malware Cloud Lookup** or **Block Malware** action and the Cloud-Delivered Firewall Management Center cannot establish connectivity with the AMP cloud, the system cannot perform any configured rule action options until connectivity is restored.

General file rule best practices

- A rule configured to block files in a passive deployment does not block matching files. Because the connection continues to transmit the file, if you configure the rule to log the beginning of the connection, you may see multiple events logged for this connection.
- A policy can include multiple rules. When you create the rules, ensure that no rule is "shadowed" by a previous rule.
- Cisco recommends that you enable **Reset Connection** for the **Block Files** and **Block Malware** actions to prevent blocked application sessions from remaining open until the TCP connection resets. If you do not reset connections, the client session will remain open until the TCP connection resets itself.

Dynamic Analysis

- The file types supported for dynamic analysis are a subset of the file types supported for other types of analysis. To view the file types supported for each type of analysis, navigate to the file rule configuration page, select the **Block Malware** action, and select the checkboxes of interest.

  To ensure that the system examines all file types, create separate rules (within the same policy) for dynamic analysis and for other types of analysis.
- If you are monitoring high volumes of traffic, do **not** store all captured files, or submit all captured files for dynamic analysis. Doing so can negatively impact system performance.

File detection best practices

Consider these notes and limitations when implementing file detection:

Prerequisite

If adaptive profiling is not enabled, access control rules cannot perform file control, including AMP.

General file detection best practices

- If a file matches a rule with an application protocol condition, file event generation occurs after the system successfully identifies a file's application protocol. Unidentified files do not generate file events.
- To detect ISO files, set the "Limit the number of bytes inspected when doing file type detection" option to a value greater than 36870, as described in File and malware inspection performance and storage options, on page 2613.
- .Exe files inside some .rar archives cannot be detected, including possibly rar5.

- If the file disposition is Neutral, it is an unknown disposition.

FTP-Specific considerations

- FTP transfers commands and data over different channels. In a passive or inline tap mode deployment, the traffic from an FTP data session and its control session may not be load-balanced to the same internal resource.

Email protocols considerations

- If the total number of bytes for all file names for files in a POP3, POP, SMTP, or IMAP session exceeds 1024, file events from the session may not reflect the correct file names for files that were detected after the file name buffer filled.
- When transmitting text-based files over SMTP, some mail clients convert newlines to the CRLF newline character standard. Since Mac-based hosts use the carriage return (CR) character and Unix/Linux-based hosts use the line feed (LF) character, newline conversion by the mail client can modify the size of the file. Note that some mail clients default to newline conversion when processing an unrecognizable file type.

File blocking best practices

Consider these notes and limitations for file blocking:

HTTP transfers

- File rules with **Block Files** and **Block Malware** actions block automatic resumption of file download via HTTP by blocking new sessions with the same file, URL, server, and client application detected for 24 hours after the initial file transfer attempt occurs.
- In rare cases, if traffic from an HTTP upload session is out of order, the system cannot reassemble the traffic correctly and therefore will not block it or generate a file event.

End-of-File detection

- If an end-of-file marker is not detected for a file, regardless of transfer protocol, the file will not be blocked by a **Block Malware** rule or the custom detection list. The system waits to block the file until the entire file has been received, as indicated by the end-of-file marker, and blocks the file after the marker is detected.
- If the end-of-file marker for an FTP file transfer is transmitted separately from the final data segment, the marker will be blocked and the FTP client will indicate that the file transfer failed, but the file will actually completely transfer to disk.

SMB or NetBIOS-ssn transfers

- If you transfer a file over NetBIOS-ssn (such as an SMB file transfer) that is blocked with a **Block Files** rule, you may see a file on the destination host. However, the file is unusable because it is blocked after the download starts, resulting in an incomplete file transfer.

- If you create file rules to detect or block files transferred over NetBIOS-ssn (such as an SMB file transfer), the system does not inspect the ongoing file transfers. However, the system inspects the new files that are transferred after you deploy an access control policy invoking the file policy.
- SMB has a functionality called multi-channel which creates multiple parallel sessions with the same IP address and different ports. For a given transaction over multi-channel, the file download is multiplexed across these sessions which is not inspected by the system as a single file.
- Files transferred concurrently in a single TCP or SMB session are not inspected.
- In a cluster environment, if an existing SMB session is moved to a new device due to a cluster role change or a device failure, then the files in any ongoing file transfers may not be inspected.
- Some SMB file transfers between Microsoft Windows systems use very high TCP window size for quick file transfers. To detect or block such file transfers, it is recommended that you increase the value of **Maximum Queued Bytes** and **Maximum Queued Segments** under **Network Analysis Policy** > **TCP Stream Configuration** > **Troubleshooting Options**.

High availability

If you configure Firewall Threat Defense high availability, and failover occurs while the original active device is identifying the file, the file type is not synced. Even if your file policy blocks that file type, the new active device downloads the file.

File policy best practices

Follow these best practices for file policy configuration to ensure proper functionality and optimal performance.

Prerequisites

- For file blocking to work, the NAP policy you apply to the access control policy must be operating in Protection mode, also known as Inline mode.
- You can associate a single file policy with an access control rule whose action is **Allow**, **Interactive Block**, or **Interactive Block with reset**.
- You **cannot** use a file policy to inspect traffic handled by the access control default action.

Encrypted traffic

By default, file inspection of encrypted payloads is disabled. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has file inspection configured.



Attention The File Inpsect preprocessor with the following generator IDs (GIDs) are enabled by default for file/malware policy: GID: 146 and GID: 147.

General file policy best practices

- For a new policy, the web interface indicates that the policy is not in use. If you are editing an in-use file policy, the web interface tells you how many access control policies use the file policy. In either case, you can click the text to jump to the Access Control Policies page.

- Based on your configuration, you can either inspect a file the first time the system detects it, and wait for a cloud lookup result, or pass the file on this first detection without waiting for the cloud lookup result.
- If you enable an access control policy with a file policy that uses either a Malware action or a Store Files option, the computing power and system performance of the device is reduced.
- NEW_OFFICE file type in file policy configuration identifies Microsoft files: DOCX, XLSX, and PPTX. When you apply a policy action to NEW_OFFICE, the action applies uniformly to all three file types—DOCX, XLSX, and PPTX.

  If your security requirements need differentiated handling of Office file formats, such as blocking DOCX documents while permitting XLSX spreadsheets, do not use the NEW_OFFICE aggregate file type. Instead, perform these steps:

  1. Select individual file types (DOCX, XLSX, PPTX) separately in your file policy rule.
  2. Create separate rules with distinct actions for each format you need to control granularly.
  3. Test your policy to ensure each Office format is handled according to your security intent.

  This approach ensures that your file policy aligns with your organization's security posture and prevents unintended blocking of legitimate business traffic.

Configure malware protection

You need to set up your system to protect your network from malicious software by following a series of configuration steps.

Procedure

Step 1 Plan and prepare for malware protection, on page 2590

Step 2 Configure file policies, on page 2591

Step 3 Add file policies to your access control configuration, on page 2591

Step 4 Configure network discovery policies to associate file and malware events with hosts on your network.

(Do not simply turn on network discovery; you must configure it to discover hosts on your network to build a network map of your organization.)

See Network Discovery Policies and subtopics.

Step 5 Deploy policies to managed devices.

See Deploy configuration changes, on page 202.

Step 6 Test your system to be sure it is processing malicious files as you expect it to.

Step 7 Verify and monitor malware protection, on page 2593

What to do next

- (Optional) To further enhance detection of malware in your network, deploy and integrate Cisco's Secure Endpoint product. Refer to (Optional) Malware protection with Secure Endpoint, on page 2615 and subtopics.

Plan and prepare for malware protection

This procedure is the first set of steps in the complete process for configuring your system to provide malware protection.

Procedure

Step 1 Purchase and install licenses.

Refer to License requirements for file and malware policies, on page 2585 and *Licenses* in the *Cisco Secure Firewall Management Center Administration Guide*.

Step 2 Understand how file policies and malware protection fit into your access control plan.

Refer to the chapter Access Control Policies, on page 1997.

Step 3 Understand the file analysis and malware protection tools.

Refer to File rule actions, on page 2604 and subtopics.

Consider also Advanced and archive file inspection options, on page 2598.

Step 4 Determine whether you will use public clouds or private (on-premises) clouds for malware protection (file analysis and dynamic analysis.)

Refer to Cloud connections for malware protection, on page 2594 and subtopics.

Step 5 If you will use private (on-premises) clouds for malware protection: Purchase, deploy, and test those products.

For information, contact your Cisco sales representative or authorized reseller.

Step 6 Configure your firewall to allow communications with your chosen clouds.

Refer to *Security, Internet Access, and Communication Ports* in the *Cisco Secure Firewall Management Center Administration Guide*.

Step 7 Configure connections between Firepower and the malware protection clouds (public or private, as applicable).

What to do next

Continue with the next step in the malware protection workflow:

Refer to Configure malware protection, on page 2589.

Configure file policies

File policies are a critical component of the malware protection workflow that allow you to define specific rules and advanced options for file inspection and handling.

Before you begin

Complete the tasks up to this point in the malware protection workflow:

Refer to Configure malware protection, on page 2589.

Procedure

Step 1 Review file policy and file rule restrictions.

Refer to Best practices for file policies and malware detection, on page 2585 and subtopics.

Step 2 Create a file policy.

Refer to Create or edit a file policy, on page 2597.

Step 3 Create rules within your file policy.

Refer to File rules, on page 2602 and subtopics.

Step 4 Configure advanced options.

Refer to Advanced and archive file inspection options, on page 2598.

What to do next

Continue with the next step in the malware protection workflow:

Refer to Configure malware protection, on page 2589.

Add file policies to your access control configuration

An access control policy can have multiple access control rules associated with file policies. You can configure file inspection for any Allow or Interactive Block access control rule, which permits you to match different file and malware inspection profiles against different types of traffic on your network before it reaches its final destination.

Before you begin

Complete the tasks up to this point in the malware protection workflow:

See Configure malware protection, on page 2589.

Procedure

Step 1 Review guidelines for file policies in access control policies. (These are different from the file rule and file policy guidelines.)

Review File and intrusion inspection order, on page 2050.

Step 2 Associate the file policy with an access control policy.

Refer to Configure an access control rule to perform malware protection, on page 2592

Step 3 Assign the access control policy to managed devices.

See Assign access control policy to devices, on page 2014.

What to do next

Continue with the next step in the malware protection workflow:

Refer to Configure malware protection, on page 2589.

Configure an access control rule to perform malware protection

Access control rules with file policies provide comprehensive malware protection by inspecting traffic flows and applying security analysis to files traversing your network.

Caution Enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Note Inline normalization is enabled automatically when a file policy is included in an access control rule. For more information, refer to The Inline Normalization Preprocessor.

Before you begin

- Adaptive profiling **must** be enabled (its default state) for access control rules to perform file control, including AMP.
- You must be an Admin, Access Admin, or Network Admin user to perform this task.

Procedure

Step 1 In the access control rule editor (from **Policies > Access Control)**, choose an **Action** of **Allow**, **Interactive Block**, or **Interactive Block with reset**.

Step 2 Choose a **File Policy** to inspect traffic that matches the access control rule, or choose **None** to disable file inspection for matching traffic.

Step 3 (Optional) Disable logging of file or malware events for matching connections by clicking **Logging** and unchecking **Log Files**.

Note
Cisco recommends that you leave file and malware event logging enabled.

Step 4 Save the rule.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Verify and monitor malware protection

Ongoing maintenance is essential for protecting your network.

Before you begin

Configure your system to protect your network from malware.

Refer to Configure malware protection, on page 2589 and referenced procedures.

Procedure

Step 1 Ensure that your system always has the most current and effective protection.

Refer to Maintain your system: file types eligible for dynamic analysis, on page 2596.

Step 2 Configure alerts for malware-related events and health monitoring.

See the *Cisco Secure Firewall Management Center Administration Guide* for information on *Configuring malware defense Alerting* and for information about the following modules:

- Local Malware Analysis
- Security Intelligence
- Threat Data Updates on Devices
- Intrusion and File Event Rate
- AMP for Firepower Status

- Secure Endpoint Status

What to do next

Review "What to do next items" in the malware protection workflow:

Refer to Configure malware protection, on page 2589.

Cloud connections for malware protection

Connections to public or private clouds are required in order to protect your network from malware.

AMP clouds

The Advanced Malware Protection (AMP) cloud is a Cisco-hosted server that uses big data analytics and continuous analysis to provide intelligence that the system uses to detect and block malware on your network.

The AMP cloud provides dispositions FOR possible malware detected IN network traffic by managed devices, as well as data updates FOR local malware analysis and file pre-classification.

If your organization has deployed Secure Endpoint and configured Firepower to import its data, the system imports this data from the AMP cloud, including scan records, malware detections, quarantines, and indications of compromise (IOC).

Cisco offers these options for obtaining data from the Cisco cloud about known malware threats:

- **AMP public cloud**: Your Cloud-Delivered Firewall Management Center communicates directly with the public Cisco cloud. There are three public AMP clouds, in the United States, Europe, and Asia.

Dynamic analysis cloud

Dynamic analysis cloud options include:

- Secure Malware Analytics Cloud: Public cloud that processes eligible files that you send for dynamic analysis, and provides threat scores and dynamic analysis reports. Firepower supports 200 samples/day FOR Secure Malware Analytics analysis.

AMP cloud connection configurations

Change AMP options

Use this procedure to modify AMP (Advanced Malware Protection) for Networks configuration settings that control automatic signature updates and data sharing with the Cisco cloud.

Procedure

Step 1 Choose **Integration** > **Other Integrations** > **Cloud Services**.

Step 2 Configure these parameters:

Table 138: AMP for networks options

Option	Description
Enable Automatic Local Malware Detection Updates	The local malware detection engine statically analyzes and preclassifies files using signatures provided by Cisco. If you enable this option, the Cloud-Delivered Firewall Management Center checks for signature updates once every 30 minutes.
Share URI from Malware Events with Cisco	The system can send information about the files detected in network traffic to the AMP cloud. This information includes URI information associated with detected files and their SHA-256 hash values. Although sharing is opt-in, transmitting this information to Cisco helps future efforts to identify and track malware.

Step 3 Click **Save**.

Dynamic analysis connections

Dynamic analysis connections are network connections that enable real-time security analysis and threat detection capabilities.

Prerequisites for dynamic analysis

This topic lists the prerequisites to use dynamic analysis.

- You must be an Admin, Access Admin, or Network Admin user, and be in the global domain, to use dynamic analysis.
- With the appropriate license, the system automatically has access to the Secure Malware Analytics Cloud.
- Dynamic analysis requires that managed devices have direct or proxied access to the Secure Malware Analytics Cloud or an on-premises Secure Malware Analytics Appliance on port 443.
- See also Files eligible for dynamic analysis, on page 2609.
-

View the default dynamic analysis connection

By default, the Cloud-Delivered Firewall Management Center can connect to the public Secure Malware Analytics Cloud for file submission and report retrieval. You can neither configure nor delete this connection.

Procedure

Step 1 Choose **Integration** > **AMP** > **Dynamic Analysis ConnectionsAMP** > **Dynamic Analysis Connections**.

Step 2 You can view the cloud in use on the default dynamic analysis connection. To associate the device, click **Associate** ( ). For more information, refer to Enable access to dynamic analysis results in the public cloud, on page 2596.

Enable access to dynamic analysis results in the public cloud

Secure Malware Analytics offers more detailed reporting on analyzed files than is available in the Cloud-Delivered Firewall Management Center. If your organization has a Secure Malware Analytics Cloud account, you can access the Secure Malware Analytics portal directly to view additional details about files sent for analysis from your managed devices. However, for privacy reasons, file analysis details are available only to the organization that submitted the files. Therefore, before you can view this information, you must associate your Cloud-Delivered Firewall Management Center with the files submitted by its managed devices.

Before you begin

You must have a Secure Malware Analytics Cloud account, and have your account credentials ready.

Procedure

Step 1 Select **Integration** > **AMP** > **Dynamic Analysis ConnectionsAMP** > **Dynamic Analysis Connections**.

Step 2 Click **Associate** ( ) in the table row corresponding to the Secure Malware Analytics Cloud.

A Secure Malware Analytics portal window opens.

Step 3 Sign in to the Secure Malware Analytics Cloud.

Step 4 Click **Submit Query**.

Note
Do not change the default value in the **Devices** field.

If you have difficulties with this process, contact your Secure Malware Analytics representative at Cisco TAC.

It may take up to 24 hours for this change to take effect.

What to do next

After the association is activated, refer to *Viewing Dynamic Analysis Results in the Cisco Cloud* in the *Cisco Secure Firewall Management Center Administration Guide*.

Maintain your system: file types eligible for dynamic analysis

The list of file types eligible for Dynamic Analysis is determined by the vulnerability database (VDB), which is updated periodically (but no more than once per day.) If you are an Admin user, you can update file types eligible for dynamic analysis.

The supported file types for automatic submission can vary in every release. The currently supported files for automatic submission are:

- PE Executable (32-bit Only)

- EXE
- DLL
- PDF
- MSOLE2 (Microsoft Object Linking and Embedding Compound File)
- DOCX, PPTX, XLSX

Note GZ and ZIP file types are not supported for automatic submission.

Procedure

Step 1 Do one of the following:

- (Recommended) Refer to *Vulnerability Database Update Automation* as discussed in the *Cisco Secure Firewall Management Center Administration Guide*
- Regularly check for new VDB updates, and *Manually Update the VDB* as discussed in the *Cisco Secure Firewall Management Center Administration Guide* when needed.

  If you choose this option, we recommend that you schedule regular reminders to do this.

Step 2 If your file policies specify individual file types instead of the **Dynamic Analysis Capable** file type category, update your file policies to use the newly supported file types.

Step 3 If the list of eligible file types changes, deploy to managed devices.

File policies and file rules

File policies and file rules work together to control file handling and inspection in network traffic.

Create or edit a file policy

File policies allow you to control file transfers and provide malware protection by defining rules that specify which files can be transferred, blocked, or inspected for threats.

Before you begin

If you are configuring policies for malware protection, refer to all required procedures in Configure file policies, on page 2591.

Procedure

Step 1 Select **Policies** > **Access Control heading** > **Malware & File** .

Step 2 Create a new policy, or edit an existing policy.

If you are editing an existing policy: If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Tip

To make a copy of an existing file policy, click **Copy** ( ), then type a unique name for the new policy in the dialog box that appears. You can then modify the copy.

Step 3 Add one or more rules to the file policy as described in Create file rules, on page 2611.

Step 4 Optionally, select Advanced and configure advanced options as described in Advanced and archive file inspection options, on page 2598.

Step 5 Save the file policy.

What to do next

- If you are configuring policies for malware protection, see other required procedures in Configure file policies, on page 2591.
- Otherwise:
  - Add the file policy to an access control rule as described in Add file policies to your access control configuration, on page 2591.
  - Deploy configuration changes.

Advanced and archive file inspection options

The Advanced Settings in the file policy editor provides general options for file analysis and archive file inspection to enhance security policy configuration.

The Advanced Settings in the file policy editor has these general options:

- **First Time File Analysis**—Select this option to analyze first-seen files while AMP cloud disposition is pending. The file must match a rule configured to perform a malware cloud lookup and Spero, local malware, or dynamic analysis. If you deselect this option, files detected for the first time are marked with an Unknown disposition
- **Enable Custom Detection List**—Block files on the custom detection list.
- **Enable Clean List**—If enabled, this policy will allow files that are on the clean list.
- **If AMP Cloud disposition is Unknown, override disposition based upon threat score**—Select an option:
  - If you select **Disabled,** the system will not override the disposition provided by the AMP Cloud.

- If you set a threshold threat score, files with an AMP cloud verdict of Unknown are considered malware if their Dynamic Analysis score is equal to or worse than the threshold.
- If you select a lower threshold value, you increase the number of files treated as malware. Depending on the action selected in your file policy, this can result in an increase of blocked files.

The Advanced Settings in the file policy editor has these archive file inspection options:

- **Inspect Archives**—Enables inspection of the contents of archive files, for archive files as large as the **Maximum file size to store** advanced access control setting.
- **Block Encrypted Archives**—Blocks password-protected archives.
- **Block Uninspectable Archives**—Blocks archive files with contents that the system is unable to inspect for reasons other than encryption. This usually applies to corrupted files, or those that exceed your specified maximum archive depth.
- **Max Archive Depth**—Blocks nested archive files that exceed the specified depth. The top-level archive file is not considered in this count; depth begins at 1 with the first nested file.

Archive files

Archive files are files that contain other files, such as `.zip` or `.rar` files.

If any individual file in an archive matches a file rule with a block action, the system blocks the entire archive, not just the individual file.

Archive file inspection capabilities

The system provides various options for inspecting archive files.

- **File types**: A complete list of inspectable archive file types appears in the Cloud-Delivered Firewall Management Center web interface on the file rule configuration page. To view that page, see Create file rules, on page 2611. Contained files that can be inspected appears in the same page.
- **File size**: You can inspect archive files as large as the **Maximum file size to store** file policy advanced access control setting.
- **Nested archives**: Archive files can contain other archive files, which can in turn contain archive files. The level at which a file is nested is its *archive file depth*. Note that the top-level archive file is not included in the depth count; depth begins at 1 with the first nested file. The system can inspect up to three levels of nested files beneath the outermost archive file (level 0). You can configure your file policy to block archive files that exceed that depth (or a lower maximum depth that you specify). If you choose not to block files that exceed the maximum archive file depth of 3, when archive files that contain some extractable contents and some contents nested at a depth of 3 or greater appear in monitored traffic, the system examines and reports data only for the files it was able to inspect. All features applicable to uncompressed files (such as dynamic analysis and file storage) are available for nested files inside archive files.
- **Encrypted files**: You can configure the system to block archives whose contents are encrypted or otherwise cannot be inspected.
- **Archives that are not inspected**: If traffic that contains an archive file is on a Security Intelligence Block list or Do Not Block list, or if the top-level archive file's SHA-256 value is on the custom detection list, the system does not inspect the contents of the archive file. If a nested file is blocked, the entire archive is blocked; however, if a nested file is allowed, the archive is not automatically passed (depending

on any other nested files and characteristics). .Exe files inside some .rar archives cannot be detected, including possibly rar5.

For details about options for archive file inspection, see Advanced and archive file inspection options, on page 2598.

Archive file dispositions

Archive file dispositions are based on the dispositions assigned to the files inside the archive. All archives that contain identified malware files receive a disposition of `Malware`. Archives without identified malware files receive a disposition of `Unknown` if they contain any unknown files, and a disposition of `Clean` if they contain only clean files.

Table 139: Archive file disposition by contents

Archive File Disposition	Number of Unknown Files	Number of Clean Files	Number of Malware Files
`Unknown`	1 or more	Any	0
`Clean`	0	1 or more	0
`Malware`	Any	Any	1 or more

Archive files, like other files, may have dispositions of `Custom Detection` or `Unavailable` if the conditions for those dispositions apply.

Archive contents and details

If your file policy is configured to inspect archive file contents, you can use the context menu in a table on pages under the Analysis > Files menu, and the network file trajectory viewer to view information about the files inside an archive when the archive file appears in a file event, malware event, or as a captured file.

All file contents of the archive are listed in table form, with a short summary of their relevant information: name, SHA-256 hash value, type, category, and archive depth. A network file trajectory icon appears by each file, which you can click to view further information about that specific file.

Override file disposition using custom lists

Override file disposition using custom lists is a security mechanism that

- allows administrators to correct incorrect file dispositions from the AMP cloud by adding SHA-256 values to specific file lists
- enables treating files as clean or malware regardless of the original cloud disposition, and
- provides per-file policy control without reevaluating the file's disposition on subsequent detection.

File list types and usage

If a file has a disposition in the AMP cloud that you know to be incorrect, you can add the file's SHA-256 value to a file list that overrides the disposition from the cloud:

- To treat a file as if the AMP cloud assigned a clean disposition, add the file to the *clean list*.

- To treat a file as if the AMP cloud assigned a malware disposition, add the file to the *custom detection list*.

On subsequent detection, the device either allows or blocks the file without reevaluating the file's disposition. You can use the clean list or custom detection list per file policy.



Note To calculate a file's SHA-256 value, you must configure a rule in the file policy to either perform a malware cloud lookup or block malware on matching files.

For complete information about using file lists in Firepower, refer to File lists, on page 2973.

Alternatively, if applicable, use Centralized file lists from Secure Endpoint, on page 2601.

Centralized file lists from Secure Endpoint

If your organization has deployed Secure Endpoint, Firepower can use Block and Allow lists created in Secure Endpoint when it queries the AMP cloud for file dispositions.

Requirements for centralized file lists

Your deployment must meet these requirements:

- Your organization must be using the AMP public cloud.
- Your organization has deployed Secure Endpoint.
- You have registered your system to Secure Endpoint using the procedure in Integrate Firepower and Secure Endpoint, on page 2618.

To create and deploy these lists, refer to the documentation or online help for Secure Endpoint.



Note File lists created in Firepower override file lists created in Secure Endpoint.

Manage file policies

The **File Policies** page displays a list of existing file policies along with their last-modified dates. You can use this page to manage your file policies.



Note The system checks for updates to the list of file types eligible for dynamic analysis (no more than once a day). If the list of eligible file types changes, this constitutes a change in the file policy; any access control policy using the file policy is marked out-of-date if deployed to any devices. You must deploy policies before the updated file policy can take effect on the device. See Maintain your system: file types eligible for dynamic analysis, on page 2596.

Procedure

Step 1 Select **Policies** > **Access Control heading** > **Malware & File** .

Step 2 Manage your file policies:

- Compare—Click **Compare Policies**; see Comparing policies.
- **Create**— To create a file policy, click **New File Policy** and proceed as described in Create or edit a file policy, on page 2597.
- **Copy**— To copy a file policy, click **Copy** ( ).

  If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- **Delete**— If you want to delete a file policy, click **Delete** ( ), then click **Yes** and **OK** as prompted.

  If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Deploy—Choose **Deploy** > **Deploy**; see Deploy configuration changes, on page 202.
- **Edit**— If you want to modify an existing file policy, click **Edit** ( ).
- **Report**—Click **Report**( ); Refer to Generate current policy reports, on page 220.

File rules

A file policy, like its parent access control policy, contains rules that determine how the system handles files that match the conditions of each rule. You can configure separate file rules to take different actions for different file types, application protocols, or directions of transfer.

File rule actions

When a file matches a rule, the rule can:

- allow or block files based on simple file type matching
- block files based on disposition (whether or not evaluation indicates that it is malicious)
- store files to the device (For information, see Captured files and file storage, on page 2609)
- submit stored (captured) files for local malware, Spero, or dynamic analysis

In addition, the file policy can:

- automatically treat a file as if it is clean or malware based on entries in the clean list or custom detection list
- treat a file as if it is malware if the file's threat score exceeds a configurable threshold

- inspect the contents of archive files (such as `.zip` or `.rar`)
- block archive files whose contents are encrypted, nested beyond a specified maximum archive depth, or otherwise uninspectable

File rule components

File rules consist of several components that determine how the system detects and handles files transmitted over various application protocols. Understanding these components helps configure effective file detection and inspection policies.

Table 140: File rule components

File Rule Component	Description
application protocol	The system can detect and inspect files transmitted via FTP, HTTP, SMTP, IMAP, POP3, and NetBIOS-ssn (SMB). **Any**, the default, detects files in HTTP, SMTP, IMAP, POP3, FTP, and NetBIOS-ssn (SMB) traffic. To improve performance, you can restrict file detection to only one of those application protocols on a per-file rule basis.
direction of transfer	You can inspect incoming FTP, HTTP, IMAP, POP3, and NetBIOS-ssn (SMB) traffic for downloaded files; you can inspect outgoing FTP, HTTP, SMTP, and NetBIOS-ssn (SMB) traffic for uploaded files. **Tip** Use **Any** to detect files over multiple application protocols, regardless of whether users are sending or receiving.
file categories and types	The system can detect various types of files. These file types are grouped into basic categories, including multimedia (swf, mp3), executables (exe, torrent), and PDFs. You can configure file rules that detect individual file types, or on entire categories of file types. For example, you could block all multimedia files, or just ShockWave Flash (swf) files. Or, you could configure the system to alert you when a user downloads a BitTorrent (torrent) file. Note that executables include file types that can run macros and scripts, since these can contain malware. For a list of file types the system can inspect, select **Policies** > **Access Control heading** > **Malware & File** , create a temporary new file policy, then click **Add Rule**. Select a file type category and the file types that the system can inspect appear in the **File Types** list. **Note** Frequently triggered file rules can affect system performance. For example, detecting multimedia files in HTTP traffic (YouTube, for example, transmits significant Flash content) could generate an overwhelming number of events.

File Rule Component	Description
file rule action	A file rule's action determines how the system handles traffic that matches the conditions of the rule. Depending on the selected action, you can configure whether the system stores the file or performs Spero, local malware, or dynamic analysis on a file. If you select a Block action, you can also configure whether the system also resets the blocked connection. For descriptions of these actions and options, see File rule actions, on page 2604. File rules are evaluated in rule-action, not numerical, order. For details, see File rule actions: evaluation order, on page 2611.

File rule actions

File rules give you granular control over which file types you want to log, block, or scan for malware. Each file rule has an associated action that determines how the system handles traffic that matches the conditions of the rule. To be effective, a file policy must contain one or more rules. You can use separate rules within a file policy to take different actions for different file types, application protocols, or directions of transfer.

File rule actions include these types:

- *Detect Files* rules allow you to log the detection of specific file types to the database, while still allowing their transmission.
- *Block Files* rules allow you to block specific file types. You can configure options to reset the connection when a file transfer is blocked, and store captured files to the managed device.
- *Malware Cloud Lookup* rules allow you to obtain and log the disposition of files traversing your network, while still allowing their transmission.
- *Block Malware* rules allow you to calculate the SHA-256 hash value of specific file types, query the AMP cloud to determine if files traversing your network contain malware, then block files that represent threats.

File rule action options

Depending on the action you select, you have different options:

File Rule Action Option	Block Files capable?	Block Malware capable?	Detect Files capable?	Malware Cloud Lookup capable?
Spero Analysis* for MSEXE	no	yes, you can submit executable files	no	yes, you can submit executable files
Dynamic Analysis*	no	yes, you can submit executable files with Unknown file dispositions	no	yes, you can submit executable files with Unknown file dispositions
Capacity Handling	no	yes	no	yes

File Rule Action Option	Block Files capable?	Block Malware capable?	Detect Files capable?	Malware Cloud Lookup capable?
Local Malware Analysis*	no	yes	no	yes
Reset Connection	yes (recommended)	yes (recommended)	no	no
Store files	yes, you can store all matching file types	yes, you can store file types matching the file dispositions you select	yes, you can store all matching file types	yes, you can store file types matching the file dispositions you select

* For complete information about these options, refer to Malware protection options, on page 2605 and its subtopics.



Caution Enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Malware protection options

The system applies several methods of file inspection and analysis to determine whether a file contains malware.

File inspection sequence

Depending on the options you enable in a file rule, the system inspects files using these tools, in order:

1. Spero analysis, on page 2607 and AMP cloud lookup, on page 2607
2. Local malware analysis, on page 2608
3. Dynamic analysis, on page 2608

For a comparison of these tools, see Comparison of malware protection options, on page 2605.

(You can also, if you choose, block all files based on their file type. For more information, see Block all files by type, on page 2611.)

Refer to information about Cisco's Secure Endpoint product at (Optional) Malware protection with Secure Endpoint, on page 2615 and subtopics.

Comparison of malware protection options

This table details the benefits and drawbacks of each type of file analysis, as well as the way each malware protection method determines a file's disposition.

Analysis Type	Benefit	Limitations	Malware Identification
Spero analysis	Structural analysis of executable files, submits Spero signature to the AMP Cloud for analysis	Less thorough than local malware analysis or dynamic analysis, only for executable files	Disposition changes from Unknown to Malware only on positive identification of malware.
Local malware analysis	Consumes fewer resources than dynamic analysis, and returns results more quickly, especially if the detected malware is common	Less thorough results than dynamic analysis	Disposition changes from Unknown to Malware only on positive identification of malware.
Dynamic analysis	Thorough analysis of unknown files using Secure Malware Analytics	Eligible files are uploaded to the public cloud or an on-premises appliance. It takes some time to complete analysis	Threat score determines maliciousness of a file. Disposition can be based on the threat score threshold configured in the file policy.
Spero analysis and local malware analysis	Consumes fewer resources than configuring local malware analysis and dynamic analysis, while still using AMP cloud resources to identify malware	Less thorough than dynamic analysis, Spero analysis only for executable files	Disposition changes from Unknown to Malware only on positive identification of malware.
Spero analysis and dynamic analysis	Uses full capabilities of AMP cloud in submitting files and Spero signatures	Results obtained less quickly than if using local malware analysis	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes based on configured threat score threshold in the file policy, and from Unknown to Malware if the Spero analysis identifies malware.
Local malware analysis and dynamic analysis	Thorough results in using both types of file analysis	Consumes more resources than either alone	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes from Unknown to Malware if local malware analysis identifies malware, or based on configured threat score threshold in the file policy.

Analysis Type	Benefit	Limitations	Malware Identification
Spero analysis, local malware analysis and dynamic analysis	Most thorough results	Consumes most resources in running all three types of file analysis	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes from Unknown to Malware if Spero analysis or local malware analysis identifies malware, or based on configured threat score threshold in the file policy.
(Block transmission of all files of a specified file type)	Does not require a Malware Defense license (This option is not technically a malware protection option.)	Legitimate files will also be blocked	(No analysis is performed.)

Note Preclassification does not itself determine a file's disposition; it is merely one of the factors that determine whether a file is eligible for Dynamic Analysis.

Spero analysis

Spero analysis examines structural characteristics such as metadata and header information in executable files. After generating a Spero signature based on this information, if the file is an eligible executable file, the device submits it to the Spero heuristic engine in the AMP cloud. Based on the Spero signature, the Spero engine determines whether the file is malware. You can also configure rules to submit files for Spero analysis without also submitting them to the AMP cloud.

Note You cannot manually submit files for Spero analysis.

AMP cloud lookup

For files that are eligible for assessment using Advanced Malware Protection, the Cloud-Delivered Firewall Management Center performs a malware cloud lookup, querying the AMP cloud for the file's disposition based on its SHA-256 hash value.

Performance optimization

To improve performance, the system caches dispositions returned by the cloud and uses the cached disposition for known files rather than querying the AMP cloud. For more information about this cache, see Cached disposition longevity, on page 2608.

Local malware analysis

Local malware analysis allows a managed device to locally inspect executables, PDFs, office documents, and other types of files for the most common types of malware, using a detection rule set provided by the Talos Intelligence Group. Because local analysis does not query the AMP cloud, and does not run the file, local malware analysis saves time and system resources.

Local malware analysis behavior

When the system identifies malware through local malware analysis, it updates the existing file disposition from Unknown to Malware and generates a new malware event. If the system does not identify malware, it does not update the file disposition from Unknown to Clean.

After the system runs local malware analysis, it caches file information such as SHA-256 hash value, timestamp, and disposition, so that if detected again within a certain period of time, the system can identify malware without additional analysis. For more information about the cache, see Cached disposition longevity, on page 2608.

Local malware analysis does not require establishing communications with the Secure Malware Analytics Cloud. However, you must configure communications with the cloud to submit files for dynamic analysis, and to download updates to the local malware analysis ruleset.

Cached disposition longevity

Dispositions returned from an AMP cloud query, associated threat scores, and dispositions assigned by local malware analysis, have a time-to-live (TTL) value. After a disposition has been held for the duration specified in the TTL value without update, the system purges the cached information. Dispositions and associated threat scores have the following TTL values:

TTL values for dispositions and threat scores

Dispositions and associated threat scores have these TTL values:

- Clean — 4 hours
- Unknown — 1 hour
- Malware — 1 hour

If a query against the cache identifies a cached disposition that timed out, the system re-queries the local malware analysis database and the AMP cloud for a new disposition.

Dynamic analysis

Dynamic analysis is a security analysis method that

- automatically submits files for behavioral analysis using Secure Malware Analytics (formerly Threat Grid), Cisco's file analysis and threat intelligence platform,
- runs files in a sandbox environment to analyze behavior and determine maliciousness, and
- returns a threat score indicating the likelihood that a file contains malware.

Dynamic analysis process and capabilities

You can configure your file policy to automatically submit files for dynamic analysis using Secure Malware Analytics (formerly Threat Grid), Cisco's file analysis and threat intelligence platform.

Devices submit eligible files to Secure Malware Analytics (either the public cloud or to an on-premises appliance, whichever you have specified) regardless of whether the device stores the file.

Secure Malware Analytics runs the file in a sandbox environment, analyzes the file's behavior to determine whether the file is malicious, and returns a threat score that indicates the likelihood that a file contains malware. From the threat score, you can view a dynamic analysis summary report with the reasons for the assigned threat score. You can also look in Secure Malware Analytics to view detailed reports for files that your organization submitted, as well as scrubbed reports with limited data for files that your organization did not submit.

For more information about Cisco Secure Malware Analytics, refer to https://www.cisco.com/c/en/us/products/security/threat-grid/index.html

To configure your system to perform dynamic analysis, see the topics under Dynamic analysis connections, on page 2595.

Files eligible for dynamic analysis

A file's eligibility for dynamic analysis depends on the file type, the file size, and the file rule's action.

Additional requirements:

- The system submits only files that match the file rules you configure.
- The file must have a malware cloud lookup disposition of Unknown or Unavailable at the time the file is sent for analysis.
- The system must preclassify the file as potential malware.

Dynamic analysis and capacity handling

Capacity handling is a system feature that temporarily stores files that are otherwise eligible for dynamic analysis if the system is temporarily unable to submit files to the cloud, either because the device cannot communicate with the cloud or because the maximum number of submissions has been reached. The system submits the stored files when the hindering condition has passed.

Storage options

Some devices can store files on the device hard drive or in a malware storage pack. Refer Malware storage packs, on page 2610.

Captured files and file storage

The file storage feature allows you to capture selected files detected in traffic, and automatically store a copy of the file temporarily to a device’s hard drive, or, if installed, to the malware storage pack.

File capture and storage capabilities

After your device captures files, you can:

- Store captured files on the device's hard drive for later analysis.
- Download the stored file to a local computer for further manual analysis or archival purposes.
- Manually submit eligible captured files for AMP cloud lookup or dynamic analysis.

Note that once a device stores a file, it will not re-capture it if the file is detected in the future and the device still has that file stored.

Note When a file is detected for the first time on your network, you can generate a file event that represents the file's detection. However, if your file rule performs a malware cloud lookup, the system requires additional time to query the AMP cloud and return a disposition. Due to this delay, the system cannot store this file until the second time it is seen on your network, and the system can immediately determine the file's disposition.

Whether the system captures or stores a file, you can:

- Review information about the captured file from **Analysis** > **Unified Events**, including whether the file was stored or submitted for dynamic analysis, file disposition, and threat score, allowing you to quickly review possible malware threats detected on your network.
- View the file's trajectory to determine how it traversed your network and which hosts have a copy.
- Add the file to the clean list or custom detection list to always treat the file as if it had a clean or malware disposition on future detection.

You configure file rules in a file policy to capture and store files of a specific type, or with a particular file disposition, if available. After you associate the file policy with an access control policy and deploy it to your devices, matching files in traffic are captured and stored. You can also limit the minimum and maximum file sizes to store.

Stored files are not included in system backups.

Malware storage packs

A malware storage pack

- provides additional storage capacity for file data when installed in security devices
- allows the primary hard drive to store events and configuration files by dedicating the entire storage pack to captured files, and
- automatically receives stored files from the primary hard drive when installed on devices that already contain file data.

Storage allocation and file management

Based on your file policy configuration, your device may store a substantial amount of file data to the hard drive. You can install a malware storage pack in the device; the system stores files to the malware storage pack, allowing more room on the primary hard drive to store events and configuration files. The system periodically deletes older files. If the device's primary hard drive does not have enough available space, and does not have an installed malware storage pack, you cannot store files.

Caution Do not attempt to install a hard drive that was not supplied by Cisco in your device. Installing an unsupported hard drive may damage the device. Malware storage pack kits are available for purchase **only** from Cisco. Contact Support if you require assistance with the malware storage pack.

Without a malware storage pack installed, when you configure a device to store files, it allocates a set portion of the primary hard drive's space to captured file storage. If you configure capacity handling to temporarily

store files for dynamic analysis, the system uses the same hard drive allocation to store these files until it can resubmit them to the cloud.

When you install a malware storage pack in a device and configure file storage or capacity handling, the device allocates the entire malware storage pack for storing these files. The device cannot store any other information on the malware storage pack.

When the allocated space for captured file storage fills to capacity, the system deletes the oldest stored files until the allocated space reaches a system-defined threshold. Based on the number of files stored, you may see a substantial drop in disk usage after the system deletes files.

If a device has already stored files when you install a malware storage pack, the next time you restart the device, any captured files or capacity handling files stored on the primary hard drive are moved to the malware storage pack. Any future files the device stores are stored to the malware storage pack.

For more information on using MSP on the Firepower devices, see the Firepower Hardware Installation Guide for your device.

Block all files by type

If your organization wants to block not only the transmission of malware files, but all files of a specific type, regardless of whether the files contain malware, you can do so.

File type categories and licensing

File control is supported for all file types where the system can detect malware, plus many additional file types. These file types are grouped into basic categories, such as multimedia (swf, mp3), executables (exe, torrent), and PDFs.

Blocking all files based on their type is not technically a malware protection feature; it does not require a Malware Defense license and does not query the AMP cloud.

File rule actions: evaluation order

A file policy will likely contain multiple rules with different actions for different situations. If more than one rule can apply to a particular situation, the evaluation order described in this topic applies. In general, simple blocking takes precedence over malware inspection and blocking, which takes precedence over simple detection and logging.

The order of precedence of file-rule actions is:

- *Block Files*
- *Block Malware*
- *Malware Cloud Lookup*
- *Detect Files*

Create file rules

Use this procedure when you need to add file rules to an existing file policy to control file handling and security actions.

Caution Enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Before you begin

If you are configuring rules for malware protection, refer to Configure file policies, on page 2591.

Procedure

Step 1 Select **Policies** > **Access Control heading** > **Malware & File** .

Step 2 Click the edit icon to modify an existing file policy.

Step 3 In the file policy editor, click **Add Rule**.

Step 4 Select an **Application Protocol** and **Direction of Transfer** as described in File rule components, on page 2603.

Step 5 Select one or more **File Types**.

The file types you see depend on the selected application protocol, direction of transfer, and action.

You can filter the list of file types in these ways:

- Select one or more **File Type Categories**, then click **All types in selected Categories.**
- Search for a file type by its name or description. For example, type `Windows` in the **Search name and description** field to display a list of Microsoft Windows-specific files.

Tip
Hover your pointer over a file type to view its description.

Step 6 Select a file rule **Action** as described in File rule actions, on page 2604, with consideration for File rule actions: evaluation order, on page 2611.

The actions available to you depend on the licenses you have installed. Refer to License requirements for file and malware policies, on page 2585.

Step 7 Depending on the action you selected, configure options:

- reset the connection after blocking the file
- store files that match the rule
- enable Spero analysis*
- enable local malware analysis*
- enable dynamic analysis* and capacity handling

* For information about these options, refer to File rule actions, on page 2604 and Malware protection options, on page 2605 and its subtopics.

Step 8 Click **Add.**

Step 9 Click **Save** to save the policy.

What to do next

- If you are configuring policies for malware protection, return to Configure file policies, on page 2591.
- Deploy configuration changes.

Access control rule logging for malware protection

When the system detects a prohibited file (including malware) according to the settings in the file policy, it automatically logs an event to the Cloud-Delivered Firewall Management Center database. If you do not want to log file or malware events, you can disable this logging on a per-access-control-rule basis.

The system also logs the end of the associated connection to the Cloud-Delivered Firewall Management Center database, regardless of the logging configuration of the invoking access control rule.

Retrospective disposition changes

File dispositions can change. For example, as new information is discovered, the AMP cloud can determine that a file that was previously thought to be clean is now identified as malware, or the reverse—that a malware-identified file is actually clean. When the disposition changes for a file you queried in the past week, the AMP cloud notifies the system so it can automatically take action the next time it detects that file being transmitted. A changed disposition is called a retrospective disposition.

File and malware inspection performance and storage options

Increasing the file sizes can affect the performance of the system.

Table 141: Advanced access control file and malware defense options

Field	Description	Guidelines and Restrictions
Limit the number of bytes inspected when doing file type detection	Specifies the number of bytes inspected when performing file type detection.	0 - 4294967295 (4GB) 0 removes the restriction. The default value is the maximum segment size of a TCP packet (1460 bytes). In most cases, the system can identify common file types using the first packet. To detect ISO files, enter a value greater than 36870.

Field	Description	Guidelines and Restrictions
Allow file if cloud lookup for Block Malware takes longer than (seconds)	Specifies how long the system will hold the last byte of a file that matches a **Block Malware** rule and that does not have a cached disposition, while malware cloud lookup occurs. If the time elapses without the system obtaining a disposition, the file passes. Dispositions of Unavailable are not cached.	0 - 30 seconds Do *not* set this option to 0 without contacting Support. Cisco recommends that you use the default value to avoid blocking traffic because of connection failures.
Do not calculate SHA-256 hash values for files larger than (in bytes)	Prevents the system from storing files larger than a certain size, performing a malware cloud lookup on the files, or blocking the files if added to the custom detection list.	0 - 4294967295 (4GB) 0 removes the restriction. This value must be greater than or equal to **Maximum file size to store (bytes)** and **Maximum file size for dynamic analysis testing (bytes)**.

Maximum file size for advanced file inspection and storage (bytes)	These settings specify:
	• The file size that the system can inspect using these detectors:
	  • Spero analysis
	  • Sandboxing and preclassification
	  • Local malware analysis/ClamAV
	  • Archive inspection
	• The file size that the system can store using a file rule.
		0 - 10485760 (10MB)
		0 disables file storage.
		Must be less than or equal to **Maximum file size to store (bytes)** and **Do not calculate SHA-256 hash values for files larger than (in bytes)**.

Minimum file size for advanced file inspection and storage (bytes)		0 - 10485760 (10MB) 0 disables file storage. Must be greater than or equal to **Minimum file size to store (bytes)**, and less than or equal to **Do not calculate SHA-256 hash values for files larger than (in bytes)**.
Maximum file size for dynamic analysis testing (bytes)	Specifies the minimum file size the system can submit to the AMP cloud for dynamic analysis.	0 -10485760 (10MB) Must be less than or equal to **Maximum file size for dynamic analysis testing (bytes)** and **Do not calculate SHA-256 hash values for files larger than (in bytes)**. The file size for dynamic analysis must be within the limits defined by the minimum and maximum settings for file analysis. The system checks the AMP cloud for updates to the minimum file size you can submit (no more than once a day). If the new minimum size is larger than your current value, your current value is updated to the new minimum, and your policy is marked out-of-date.

Field	Description	Guidelines and Restrictions
Minimum file size for dynamic analysis testing (bytes)	Specifies the maximum file size the system can submit to the AMP cloud for dynamic analysis.	0 -10485760 (10MB) Must be greater than or equal to **Minimum file size for dynamic analysis testing (bytes)**, and less than or equal to **Do not calculate SHA-256 hash values for files larger than (in bytes)**. The file size for dynamic analysis must be within the limits defined by the minimum and maximum settings for file analysis. The system checks the AMP cloud for updates to the maximum file size you can submit (no more than once a day). If the new maximum size is smaller than your current value, your current value is updated to the new maximum, and your policy is marked out-of-date.

Tune file and malware inspection performance and storage

File and malware inspection can impact system performance and storage utilization. Tuning these settings helps ensure optimal security posture while maintaining acceptable system performance.

Before you begin

You must be an Admin, Access Admin, or Network Admin user to perform this task.

Procedure

Step 1 In the access control policy editor, click **Advanced Settings**.

Step 2 Click **Edit** ( ) next to **Files and Malware Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Set any of the options described in File and malware inspection performance and storage options, on page 2613.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

(Optional) Malware protection with Secure Endpoint

Cisco's Secure Endpoint is a separate malware-protection product that can supplement malware protection provided by the system and be integrated with your Firepower deployment.

Secure Endpoint is Cisco's enterprise-class Advanced Malware Protection solution that runs as a lightweight connector on individual users' endpoints (computers and mobile devices) to discover, understand, and block advanced malware outbreaks, advanced persistent threats, and targeted attacks.

Benefits of Secure Endpoint

- Configure custom malware detection policies and profiles for your entire organization, as well as perform flash and full scans on all your users' files
- Perform malware analysis, including view heat maps, detailed file information, network file trajectory, and threat root causes
- Configure multiple aspects of outbreak control, including automatic quarantines, application blocking to stop non-quarantined executables from running, and exclusion lists
- Create custom protections, block execution of certain applications based on group policy, and create custom Allowed Applications lists
- Use the Secure Endpoint management console to help you mitigate the effect of malware. The management console provides a robust, flexible web interface where you control all aspects of your Secure Endpoint deployment and manage all phases of an outbreak.

For detailed information about Secure Endpoint, refer to:

- https://www.cisco.com/c/en/us/products/security/amp-for-endpoints/index.html.
- Online help in the Secure Endpoint management console.
- Secure Endpoint documentation available from: http://docs.amp.cisco.com.

Comparison of malware protection: Firepower vs. Secure Endpoint

This table compares the advanced malware protection differences between Firepower malware protection and Secure Endpoint across various features including detection methods, network traffic inspection, malware analysis, mitigation capabilities, and licensing requirements.

Table 142: Advanced malware protection differences by detecting product

Feature	Firepower Malware Protection (malware defense)	Secure Endpoint
File type detection and blocking method (file control)	In network traffic, using access control and file policies	Not supported
Malware detection and blocking method	In network traffic, using access control and file policies	On individual endpoints (end-user computers and mobile devices), using a connector that communicates with the AMP cloud
Network traffic inspected	Traffic passing through a managed device	None; connectors installed on endpoints directly inspect files
Malware intelligence data source	AMP cloud (public or private)	AMP cloud (public or private)
Malware detection robustness	Limited file types	All file types

Feature	Firepower Malware Protection (malware defense)	Secure Endpoint
Malware analysis choices	Cloud-Delivered Firewall Management Center-based, plus analysis in the AMP cloud	Cloud-Delivered Firewall Management Center-based, plus additional options on the Secure Endpoint management console
Malware mitigation	Malware blocking in network traffic, Cloud-Delivered Firewall Management Center-initiated remediations	Secure Endpoint-based quarantine and outbreak control options, Cloud-Delivered Firewall Management Center-initiated remediations
Events generated	File events, captured files, malware events, and retrospective malware events	Malware events
Information in malware events	Basic malware event information, plus connection data (IP address, port, and application protocol)	In-depth malware event information; no connection data
Network file trajectory	Cloud-Delivered Firewall Management Center-based	Cloud-Delivered Firewall Management Center and the Secure Endpoint management console each have a network file trajectory. Both are useful.
Required licenses or subscriptions	Licenses required to perform file control and malware defense	Secure Endpoint subscription. No license is required to bring Secure Endpoint data into Cloud-Delivered Firewall Management Center.

Integrate Secure Firewall with Secure Endpoint

If your organization has deployed Secure Endpoint, you can optionally integrate that product with your Secure Firewall deployment.

Integration with Secure Endpoint does not require a dedicated Secure Firewall license.

Benefits of integrating Firepower and Secure Endpoint

Integrating your Secure Endpoint deployment with your system offers these benefits:

- Centralized Blocked Applications and Allowed Applications lists configured in Secure Endpoint can determine verdicts for file SHAs sent from Firepower to the AMP cloud for disposition.

  Refer to Centralized file lists from Secure Endpoint, on page 2601.

- The system can import malware events detected by Secure Endpoint into Cloud-Delivered Firewall Management Center so you can manage these events along with malware events generated by the system. Imported data for these events includes scans, malware detections, quarantines, blocked executions, and cloud recalls, as well as indications of compromise (IOCs) that Cloud-Delivered Firewall Management Center displays for hosts that it monitors.

- You can view file trajectory and other details in the Secure Endpoint console.



Important If you use a Cisco AMP Private Cloud, refer to limitations at Secure Endpoint and AMP private cloud, on page 2618.

Secure Endpoint and AMP private cloud

If you configure a Cisco AMP private cloud to collect Secure Endpoint data on your network, all Secure Endpoint connectors send data to the private cloud, which forwards that data to the Cloud-Delivered Firewall Management Center. The private cloud does not share any of your endpoint data over an external connection.

Private cloud functionality and limitations

The private cloud acts as an anonymized proxy to ensure the security and privacy of your monitored network. This includes importing Secure Endpoint data. The private cloud does not share any of your endpoint data over an external connection.

The integration features that are not available with an AMP private cloud are:

- Use of Blocked Applications and Allowed Applications lists configured in Secure Endpoint. (These lists are used to block or allow files.)
- Visibility in Secure Endpoint of malware events generated from Firepower.

You can configure multiple private clouds to support the capacity you require.

Integrate Firepower and Secure Endpoint

If your organization has deployed Cisco's Secure Endpoint product, you can integrate that application with Firepower to achieve the benefits described in Benefits of integrating Firepower and Secure Endpoint, on page 2617.

When you integrate with Secure Endpoint, you must configure the Secure Endpoint connection even if you already have malware defense (AMP for Firepower) connections configured. You can configure multiple Secure Endpoint cloud connections.

Note Secure Endpoint connections that have not registered successfully do not affect malware defense.

Before you begin

- You must be an Admin user to perform this task.
- Secure Endpoint must be set up and working properly on your network.
- If you are connecting to the AMP cloud after either reimaging or restoring your Cloud-Delivered Firewall Management Center from backup, use the Secure Endpoint management console to remove the previous connection.
- You will need your Secure Endpoint credentials to log in to the Secure Endpoint console during this procedure.

Procedure

Step 1 Choose **Integration** > **AMP** > **AMP Management**.

Step 2 Click **Add AMP Cloud Connection**.

Step 3 From the **Cloud Name** drop-down list, choose the cloud you want to use.

Step 4 If you want to use this cloud for both malware defense and Secure Endpoint, select the **Use for AMP for Firepower** check box.

If you configured a different cloud to handle malware defense (AMP for Firepower) communications, you can clear this check box; if this is your only AMP cloud connection, you cannot.

Step 5 Click **Register**.

A **Spinning state** () icon indicates that a connection is pending, for example, after you configure a connection on the Cloud-Delivered Firewall Management Center, but before you authorize it using the Secure Endpoint management console. A **Denied** () icon indicates that the cloud denied the connection or the connection failed for another reason.

Step 6 Confirm that you want to continue to the Secure Endpoint management console, then log into the management console.

Step 7 Using the management console, authorize the AMP cloud to send Secure Endpoint data to Cloud-Delivered Firewall Management Center.

Step 8 If you want to restrict the data that the Cloud-Delivered Firewall Management Center receives, select specific groups within your organization for which you want to receive information.

By default, the AMP cloud sends data for all groups. To manage groups, choose **Management** > **Groups** on the Secure Endpoint management console. For detailed information, see the management console online help.

Step 9 Click **Allow** to enable the connection and start the transfer of data.

Clicking **Deny** returns you to the Cloud-Delivered Firewall Management Center, where the connection is marked as denied. If you navigate away from the Applications page on the Secure Endpoint management console, and neither deny nor allow the connection, the connection is marked as pending on the Cloud-Delivered Firewall Management Center's web interface. The health monitor does **not** alert you of a failed connection in either of these situations. If you want to connect to the AMP cloud later, delete the failed or pending connection, then recreate it.

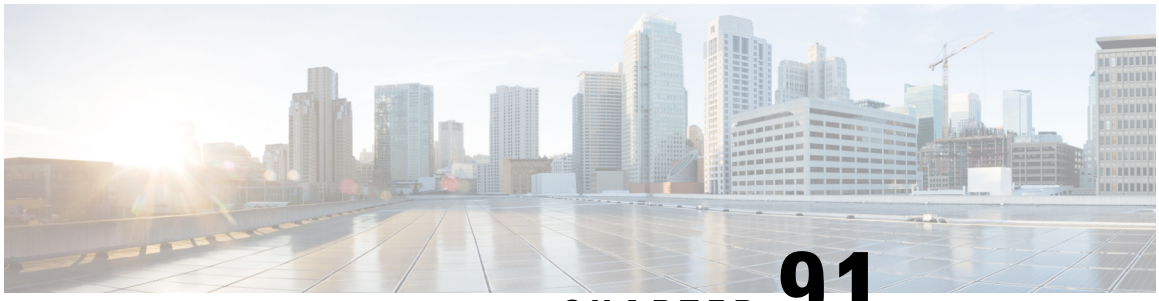
Incomplete registration of the Secure Endpoint connection does not disable the malware defense connection.

What to do next

- In the Secure Endpoint console window, configure settings as needed. For example, define group membership for your management center and assign policies. For information, see the Secure Endpoint online help or other documentation.
- The default health policy warns you if the Cloud-Delivered Firewall Management Center cannot connect to the Secure Endpoint portal after an initial successful connection, or if the connection is deregistered using the AMP portal.

  Verify that the Secure Endpoint Status monitor is enabled under **System** () > **Health** > **Policy**.
- To verify that the connection is correctly configured:
  1. On the **Integration** > **AMP** > **AMP Management** page, click the Cloud Name that includes Secure Endpoint in the **Cisco AMP Solution Type** column.
  2. In the Secure Endpoint console window that displays, choose **Accounts > Applications**.

3. Verify that your Cloud-Delivered Firewall Management Center is on the list.
4. In the Secure Endpoint console window, choose **Manage > Computers**.
5. Verify that your Cloud-Delivered Firewall Management Center is on the list.



CHAPTER 91

Dynamic Attributes Connector

The following topics discuss how to configure and use the Dynamic Attributes Connector.

- About the Dynamic Attributes Connector , on page 2621
- Enable the dynamic attributes connector, on page 2624
- Dashboard, on page 2624
- Connectors, on page 2631
- Create an adapter, on page 2659
- Create dynamic attributes filters, on page 2661
- Disable the dynamic attributes connector, on page 2664
- Dynamic firewall, on page 2664
- Use Dynamic Objects in Access Control Policies, on page 2679
- Troubleshoot the Dynamic Attributes Connector, on page 2683

About the Dynamic Attributes Connector

The dynamic attributes connector enables your access control and DNS policy to adapt in real time to the changes in public and private cloud workloads and business-critical software-as-a-service (SaaS) applications. It simplifies policy management by keeping rules up to date without tedious manual updates and policy deployment. Customers require policy rules to be defined based on non-network constructs such as VM name or security group, so that firewall policy is persistent even when the IP address or VLAN changes.

Supported connectors

We currently support:

Table 143: List of supported connectors by dynamic attributes connector version and platform

CSDAC version	AWS	AWS security groups	AWS service tags	Azure	Azure Service Tags	Cisco Cyber Vision	Cisco Multicl. Defense	Generic text	GitHub	Google Cloud	Microsoft Office 365	vCenter	Tenable	Webex	Zoom
Version 1.1 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	No	Yes	Yes	No	No	No
Version 2.0 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	Yes	Yes	Yes	No	No	No
Version 2.2 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	Yes	Yes	Yes	Yes	No	No	No

CSDAC version	AWS	AWS security groups	AWS service tags	Azure	Azure Service Tags	Cisco Cyber Vision	Cisco Multicl. Defense	Generic text	GitHub	Google Cloud	Microsoft Office 365	vCenter	Tenable	Webex	Zoom
Version 2.3 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	Yes	Yes	Yes	Yes	No	Yes	Yes
Version 3.0 (on-premises)	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Cloud-delivered (Security Cloud Control)	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes

More information about connectors:

- Amazon Web Services (AWS)

  For more information, see a resource like Tagging AWS resources on the Amazon documentation site.

  See Amazon Web Services Connector—About User Permissions and Imported Data

- Cisco Cyber Vision

  See Create a Cisco cyber vision connector, on page 2643.

- Cisco Multicl. Defense

  See Create a Multicloud Defense connector, on page 2641.

- Generic text list of IP addresses you specify

  For more information, see Create a generic text connector, on page 2645.

- GitHub

  For more information, see Create a GitHub connector, on page 2649.

- Google Cloud

  For more information, see Setting Up Your Environment in the Google Cloud documentation.

  See Google cloud connector user permissions and imported data, on page 2650.

- Microsoft Azure

  For more information, see this page on the Azure documentation site.

  See Azure Connector—About User Permissions and Imported Data.

- Microsoft Azure service tags

  For more information, see a resource like Virtual network service tags on Microsoft TechNet.

  See Create an Azure service tags connector, on page 2640

- Office 365 IP addresses

  For more information, see Office 365 URLs and IP address ranges on docs.microsoft.com.

- Tenable vulnerabilities

  For more information, see Tenable connector, on page 2654

- VMware categories and tags managed by vCenter and NSX-T

For more information, see a resource like vSphere Tags and Attributes in the VMware documentation site.

- Webex IP addresses

  For more information, see Create a Webex connector, on page 2657.

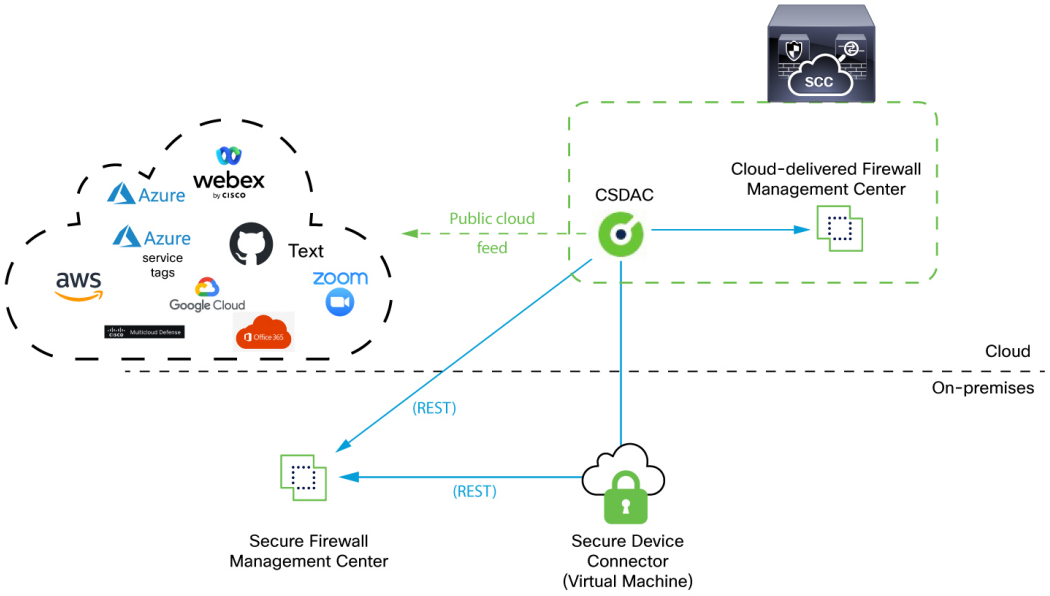
- Zoom IP addresses

  For more information, see Create a Zoom connector, on page 2658.

How It Works

This topic discusses the architecture of the dynamic attributes connector.

The following figure shows how the system functions at a high level.



- The system supports certain public cloud providers.

  This topic discusses supported *connectors* (which are the connections to those providers).

- The dynamic attributes connector is provided with Security Cloud Control; it includes a Cloud-Delivered Firewall Management Center adapter and you can connect to an On-Prem Firewall Management Center using the Secure Device Connector.

  For more information about the Secure Device Connector, see Secure Device Connector (SDC).

- The *adapter* defined by the dynamic attributes connector receives those dynamic attributes filters as *dynamic objects* and enables you to use them in access control rules.

  You can create the following types of adapters:

  - *On-Prem Firewall Management Center* for an on-premises device.

    This type of device might be managed by Security Cloud Control or it might be a standalone.

  - *Cloud-Delivered Firewall Management Center* for devices managed by Security Cloud Control.

Enable the dynamic attributes connector

Enable the dynamic attributes connector to allow objects from cloud networking products to be used in Cloud-Delivered Firewall Management Center access control and DNS rules.

This task discusses how to enable the dynamic attributes connector in the Cloud-Delivered Firewall Management Center. The dynamic attributes connector is an integration that enables objects from cloud networking products to be used in Cloud-Delivered Firewall Management Center access control and DNS rules.

Procedure

Step 1 Log in to the Cloud-Delivered Firewall Management Center if you have not done so already.

Step 2 Click **Administration** > **Dynamic Attributes Connector**.

Step 3 Slide to **Enabled**.

Step 4 The system displays messages when the dynamic attributes connector is enabled.

In the event of errors, try again. If errors persist, contact Cisco TAC.

See Connectors, on page 2631.

Dashboard

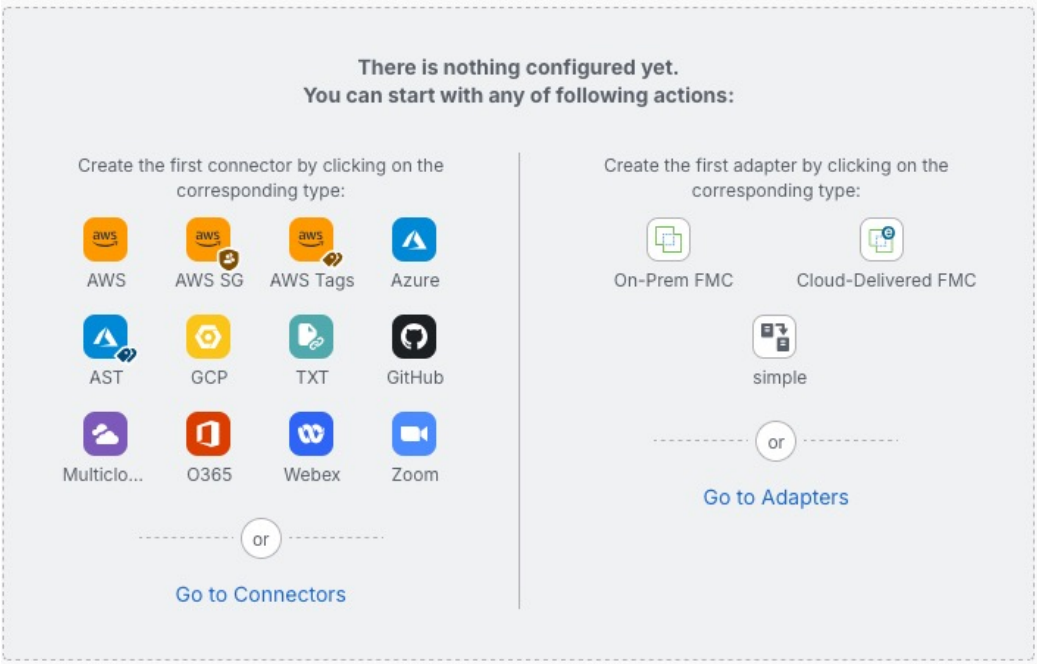
Dashboard capabilities include:

- Add, edit, and delete connectors and dynamic attributes filters.
- See how connectors and dynamic attributes filters are related to each other.
- View warnings and errors.

Dashboard access and functionality

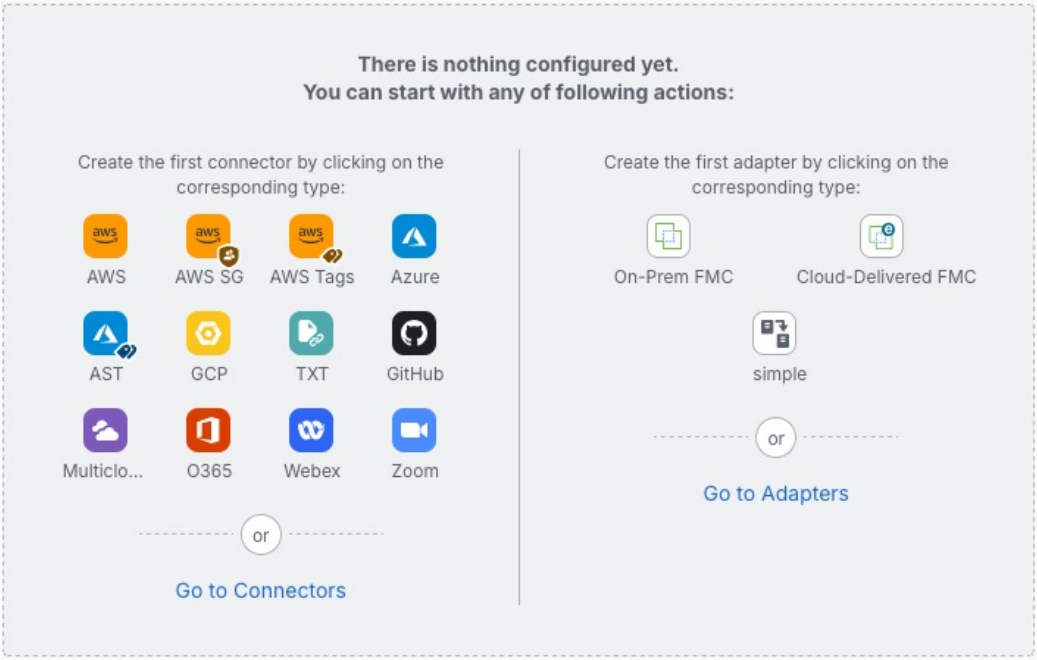
To access the dynamic attributes connector dashboard, log in to Security Cloud Control and click **Administration** > **Dynamic Attributes Connector** at the top of the page.

The dynamic attributes connector Dashboard page displays an example of an unconfigured system:

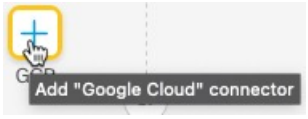


Dashboard of an unconfigured system

This sample dynamic attributes connector Dashboard page shows how an unconfigured system appears initially.



The Dashboard initially displays all the types of connectors you can configure for your system. You can do any interaction:

- Hover the mouse pointer over a connector and click 

to create a new one.
- Click **Go to Connectors** to add, edit, or delete connectors (good for creating, editing, or deleting multiple connectors at the same time).

  For more information, see Connectors, on page 2631.

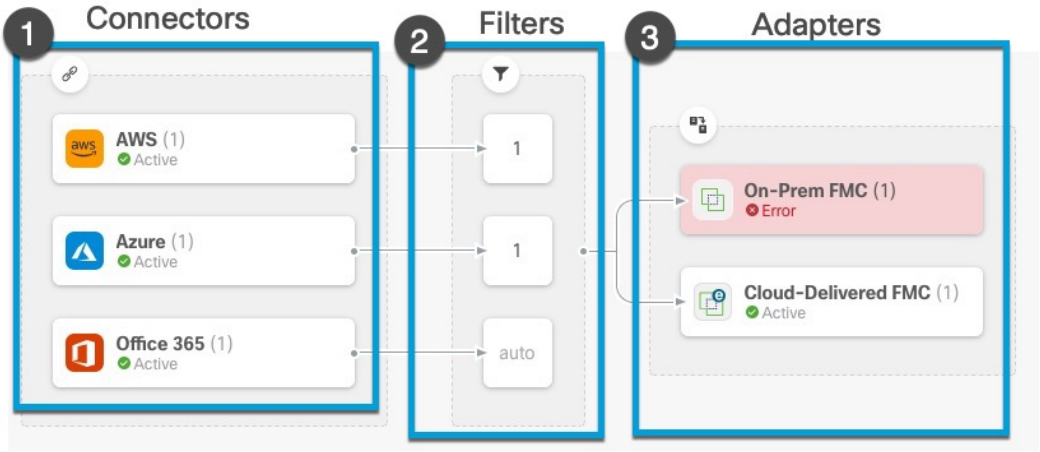
Related Topics:

- Dashboard of a configured system, on page 2626
- Add, edit, or delete connectors, on page 2627
- Add, edit, or delete dynamic attributes filters, on page 2629

Dashboard of a configured system

Sample dynamic attributes connector Dashboard page of a configured system displays at-a-glance information about configured connectors, filters, and adapters.

Click an area in the figure to learn more about it or click one of the links that follow the figure.



1 Connectors, on page 2631
2 Create dynamic attributes filters, on page 2455
3 Create an adapter, on page 2659

The Dashboard shows the following (from left to right):

Connectors column	Filters column	Adapters column
List of connectors with a number indicating how many of each type are configured. Connectors collect dynamic attributes that could be sent to the configured adapter. Dynamic attributes filters specify what data is sent. Click to view more information about all configured connectors. You can also click the name of a connector to add, edit, or delete connectors; or to view detailed information about them. For more information, see Add, edit, or delete connectors, on page 2627.	List of dynamic attributes filters associated with each connector with a number indicating how many of each filter are associated with a connector. Click to view more information about all configured filters. You can also click the name of a filter to add, edit, or delete filters; or to view detailed information about them. For more information, see Add, edit, or delete dynamic attributes filters, on page 2629.	List of adapters. Adapters receive dynamic objects from configured connectors using configured dynamic attributes filters; these dynamic objects can be used in access control policies without the need to deploy them. Click to view more information about all configured adapters. You can also click the name of an adapter to add, edit, or delete adapters; or to view detailed information about them. For more information, see Create an adapter, on page 2659.

Note Some connectors, such as Outlook 365 and Azure Service tags, automatically pull available dynamic objects without the need for a dynamic attributes filters. Those connectors display **Auto** in the column.

The Dashboard indicates whether or not an object is available. The Dashboard page is refreshed every 15 seconds but you can click **Refresh** ( ) at the top of the page at any time to refresh immediately. If issues persist, check your network connection.

Add, edit, or delete connectors

The Dashboard enables you to manage connectors by viewing, adding, editing, or deleting them. You can also view error information for connectors and troubleshoot issues.

Use the dynamic attributes connector dashboard when you need to manage connectors for your network security infrastructure. The dashboard provides comprehensive options for connector management and troubleshooting.

Procedure

Step 1 Access the connector dashboard and select your desired action.

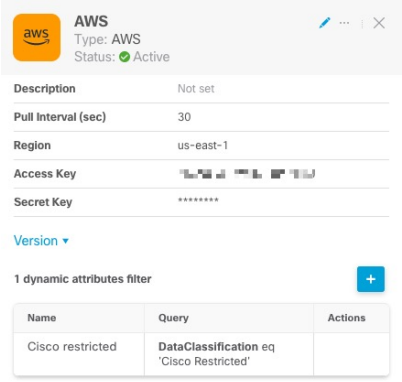
The Dashboard enables you to view or edit connectors. You can click the name of a connector to view all instances of that connector or you can click for the following additional options:

- **Go to Connectors** to view all connectors at the same time; you can add, edit, and delete connectors from there.

- **Add Connector** > ***type*** to add a connector of the indicated type.

Step 2 Click any connector in the connectors column ( ) to display more information about it.

An example follows:

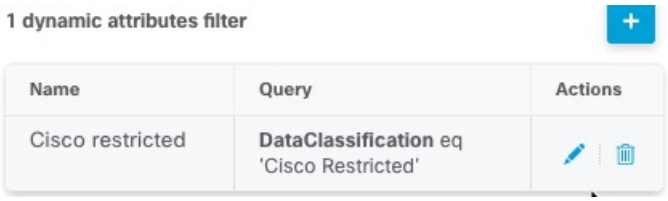


Step 3 Choose from the following management options.

- Click the Edit icon ( ) to edit this connector.
- Click the More icon ( ) for additional options.
- Click to close the panel.
- Click **Version** to display the version of the . You can optionally copy the version to the clipboard if necessary for Cisco TAC.

Step 4 Use the table at the bottom of the panel to add dynamic attributes filters or to edit or delete connectors.

A sample follows:



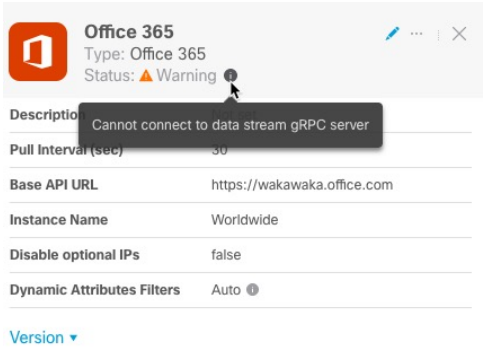
Click the Add icon ( ) to add a dynamic attributes filter for this connector. For more information, see Create dynamic attributes filters, on page 2455.

Hover the mouse pointer over the Actions column to either edit or delete the indicated connector.

Step 5 To view error information for a connector, click the name of the connector that is displaying the error.

Step 6 In the right pane, click **Information** ( ).

An example follows.



Step 7 To resolve this issue, edit the connector settings as discussed in Create an office 365 connector, on page 2653.

Step 8 If you cannot resolve the issue, click **Version** and copy the version to a text file.

Step 9 Get your Security Cloud Control tenant ID as discussed in Get your tenant ID, on page 2684.

What to do next

Provide all of this information to Cisco TAC.

Add, edit, or delete dynamic attributes filters

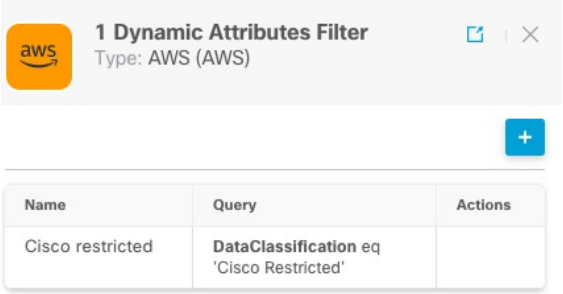
The dashboard enables you to add, edit, or delete dynamic attributes filters. You can click the name of a filter to view all instances of that filter or you can click  for additional options.

You can click the name of a filter to view all instances of that filter or you can click  for the following additional options:

- **Go to Dynamic Attributes Filters** to view all configured dynamic attributes filters. You can add, edit, or delete dynamic attributes filters from there.
- **Add Dynamic Attributes Filters** to add a filter.

For more information about adding dynamic attributes filters, see Create dynamic attributes filters, on page 2455.

An example follows:



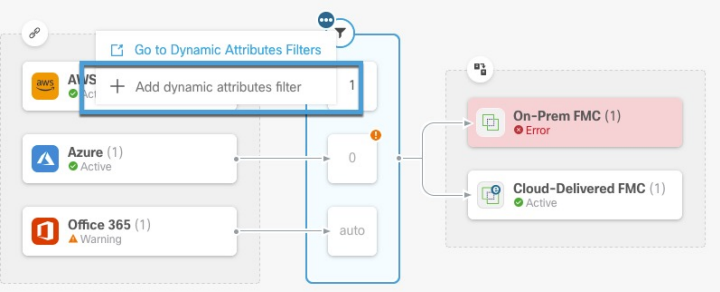
Note Some connectors, such as Outlook 365 and Azure Service tags, automatically pull available dynamic objects without the need for a dynamic attributes filters. Those connectors display **Auto** in the column.

You have these options:

- Click a filter instance to view summary information about dynamic attributes filters associated with a connector.
- Click the Add icon ( ) to add a new dynamic attributes filter.

  For more information, see Create dynamic attributes filters, on page 2455.
- Click in the filters column ( ) indicates the indicated connector has no associated dynamic attributes filters. Without associated filters, the connector can send nothing to Cloud-Delivered Firewall Management Center.

  One way to resolve the issue is to click in the filters column and click **Add Dynamic Attributes Filter**. A sample follows.



- Click to add, edit, or delete filters.
- Click to close the panel.

Connectors

A connector is an interface with a cloud service. It retrieves network information from the cloud service to enable the network information to be used in policies on the Security Cloud Control. It supports both air gap mode and network download mode configurations.

Supported connectors

We support these connectors:

Table 144: List of supported connectors by dynamic attributes connector version and platform

CSDAC version	AWS	AWS security groups	AWS service tags	Azure	Azure Service Tags	Cisco Cyber Vision	Cisco Multicl. Defense	Generic text	GitHub	Google Cloud	Microsoft Office 365	vCenter	Tenable	Webex	Zoom
Version 1.1 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	No	Yes	Yes	No	No	No
Version 2.0 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	Yes	Yes	Yes	No	No	No
Version 2.2 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	Yes	Yes	Yes	Yes	No	No	No
Version 2.3 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	Yes	Yes	Yes	Yes	No	Yes	Yes
Version 3.0 (on-premises)	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Cloud-delivered (Security Cloud Control)	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes

Amazon Web Services connector user permissions and imported data

The dynamic attributes connector imports dynamic attributes from Amazon Web Services (AWS) to Security Cloud Control for use in policies.

Dynamic attributes imported

The connector imports these dynamic attributes from AWS:

- *Tags*, user-defined key-value pairs you can use to organize your AWS EC2 resources.

  For more information, see Tag your EC2 Resources in the AWS documentation

- *IP addresses* of virtual machines in AWS.

Minimum permissions required

The dynamic attributes connector requires a user at minimum with a policy that permits `ec2:DescribeTags`, `ec2:DescribeVpcs`, and `ec2:DescribeInstances` to be able to import dynamic attributes.

Create an AWS user with minimal permissions for the dynamic attributes connector

This task creates a service account with minimum permissions to send dynamic attributes to Security Cloud Control.

For a list of the dynamic attributes that can be sent, refer to Amazon Web Services connector user permissions and imported data, on page 2631.

Before you begin

You must already have set up your Amazon Web Services (AWS) account. For more information about doing that, see this article in the AWS documentation.

Follow these steps to create an AWS user with minimal permissions for the dynamic attributes connector:

Procedure

Step 1 Log in to the AWS console as a user with the admin role.

Step 2 From the Dashboard, click **Security, Identity & Compliance** > **IAM**.

Step 3 Click **Access Management** > **Users**.

Step 4 Click **Add Users**.

a) In the **User Name** field, enter a name to identify the user.

b) Click **Access Key - Programmatic Access**.

c) At the Set permissions page, click **Next** without granting the user access to anything. You can grant user access later.

d) Add tags to the user if desired.

e) Click **Create User**.

f) Click **Download .csv** to download the user's key to your computer.

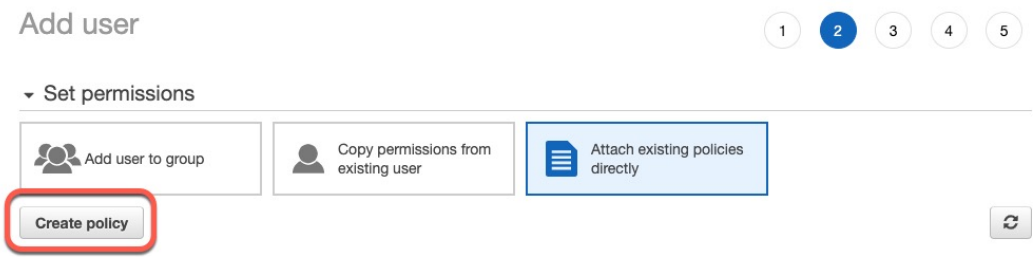
Note
This is the only opportunity you have to retrieve the user's key.

g) Click **Close**.

Step 5 At the Identity and Access Management (IAM) page in the left column, click **Access Management** > **Policies**.

Step 6 Click **Create Policy**.

a) On the Create Policy page, click **JSON**.



b) Enter the following policy in the field:

```
{
 "Version": "2012-10-17",
 "Statement": [
```

```
    {
     "Effect": "Allow",
     "Action": [
      "ec2:DescribeTags",
      "ec2:DescribeInstances",
      "ec2:DescribeVpcs"
     ],
     "Resource": "*"
    }
   ]
}
```

c) Click **Next**.

Step 7 Click **Review**.

On the Review Policy page, enter the requested information and click **Create Policy**.

Step 8 Attach your policy

a) On the Policies page, enter all or part of the policy name in the search field and press Enter.
b) Click the policy you just created.
c) Click **Actions** > **Attach**.
d) If necessary, enter all or part of the user name in the search field and press Enter.
e) Click **Attach Policy**.

What to do next

Create an AWS connector, on page 2633.

Create an AWS connector

Create a connector to send dynamic objects from AWS to the firewall for use in security policies.

This task discusses how to configure a connector that sends data from AWS to the Security Cloud Control for use in policies.

Before you begin

Create a user with at least the privileges discussed in Create an AWS user with minimal permissions for the dynamic attributes connector, on page 2632.

Follow these steps to create an AWS connector:

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of these actions:

- Add a new connector: click Add icon (), then click the name of the connector.

- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter the required information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from AWS.
Region	(Required.) Enter your AWS region code.
Access Key	(Required.) Enter your access key.
Secret Key	(Required.) Enter your secret key.

a) Click **Test** and make sure the test succeeds before you save the connector.
b) Click **Save**.

The Status column displays **Ok**.

Amazon Web Services security groups connector—About user permissions

The dynamic attributes connector imports dynamic attributes from AWS to Security Cloud Control for use in policies.

Minimum permissions required

The dynamic attributes connector requires a user at minimum with a policy that permits `ec2:DescribeTags`, `ec2:DescribeVpcs`, and `ec2:DescribeInstances` to be able to import dynamic attributes.

Create an AWS security groups connector

This task creates a connector that sends AWS security groups data to the Security Cloud Control for use in policies.

This task discusses how to configure a connector that sends AWS security groups data to the Security Cloud Control for use in policies.

Before you begin

Complete these steps:

- Create AWS security groups as discussed in Work with security groups on the AWS documentation site.
- Create a user with at least the privileges discussed in Create an AWS user with minimal permissions for the dynamic attributes connector, on page 2632.

Follow these steps to create an AWS Security Groups connector:

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from AWS. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Region	(Required.) Enter your AWS region code.
AWS Access Key	(Required.) Enter your access key.
AWS Secret Key	(Required.) Enter your secret key.

a) Click **Test** and make sure the test succeeds before you save the connector.

b) Click **Save**.

The Status column displays **Ok**.

Create an AWS service tags connector

Create a connector for Amazon Web Services (AWS) service tags to the Security Cloud Control for use in policies.

For more information, see resources like the following on the AWS documentation site:

- What are tags?
- AWS IP address ranges
- Tagging your AWS resources
- Guidance for Tagging on AWS
- AWS service points

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon (), then click the name of the connector.
- Edit a connector: click Edit icon ().
- Delete a connector: click Delete icon ().

Step 5 Enter the necessary information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
URL	(Required.) Do not change the URL unless advised to do so.

a) Click **Test** and make sure **Test connection succeeded** is displayed before you save the connector.

b) Click **Save**.

The Status column displays **Ok**.

Azure connector—About user permissions and imported data

The dynamic attributes connector imports dynamic attributes from Azure to Security Cloud Control for use in policies.

Dynamic attributes imported

We import these dynamic attributes from Azure:

- *Tags*, key-value pairs associated with resources, resource groups, and subscriptions.

For more information, see this page in the Microsoft documentation.

- *IP addresses* of virtual machines in Azure.

Minimum permissions required

The dynamic attributes connector requires a user at minimum with the **Reader** permission to be able to import dynamic attributes.

Create an Azure user with minimal permissions for the dynamic attributes connector

Create an Azure service account with minimal permissions to send dynamic attributes to Security Cloud Control.

This task discusses how to set up a service account with minimum permissions to send dynamic attributes to Security Cloud Control. For a list of these attributes, see Azure connector—About user permissions and imported data, on page 2636.

Before you begin

You must already have a Microsoft Azure account. To set one up, see this page on the Azure documentation site.

Follow these steps to create an Azure user with minimal permissions for the dynamic attributes connector:

Procedure

Step 1 Log in to the Azure Portal as the owner of the subscription.

Step 2 Click **Azure Active Directory**.

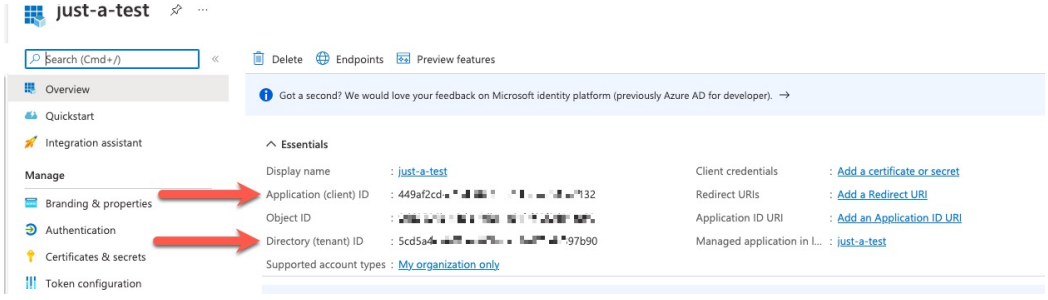
Find the instance of Azure Active Directory for the application you want to set up.

Step 3 Click **Add** > **App registration**.

a) In the **Name** field, enter a name to identify this application.
b) Enter other information on this page as required by your organization.
c) Click **Register**.

Step 4 On the next page, write down or copy the Client ID (also referred to as *application ID*) and the tenant ID (also referred to as the *directory ID*).

Here is a sample:

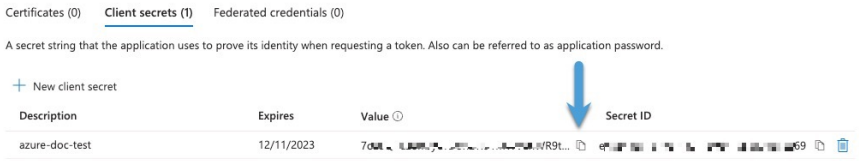


Step 5 Next to Client Credentials, click **Add a certificate or secret**.

Step 6 Click **New Client Secret**.

a) Enter the requested information and click **Add**.

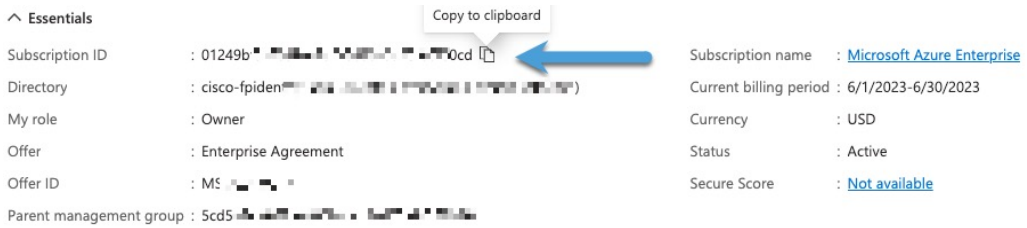
b) Copy the value of the **Value** field to the clipboard. This value, *and not the* ***Secret ID***, is the client secret.



Step 7 Go back to the main Azure Portal page and click **Subscriptions**.

a) Click the name of your subscription.

b) Copy the subscription ID to the clipboard.



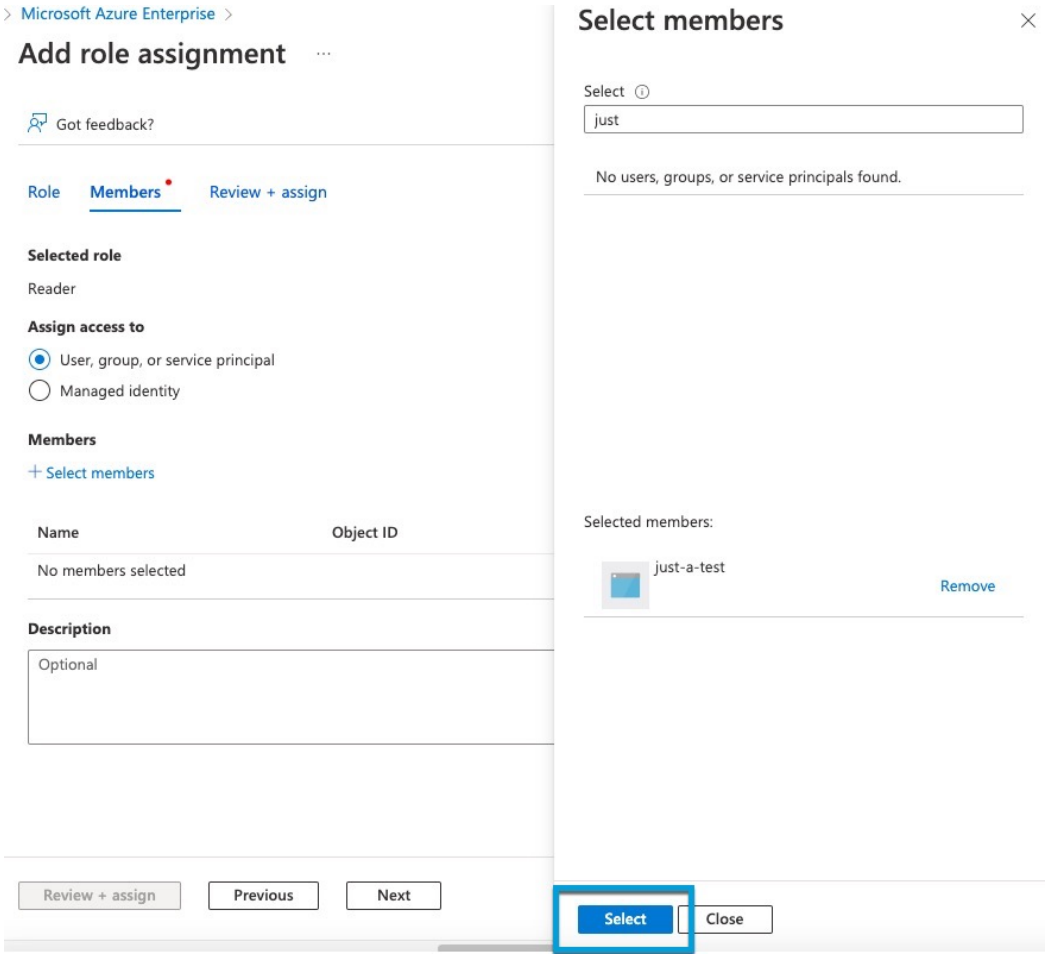
Step 8 Click **Access Control (IAM)**.

a) Click **Add** > **Add role assignment**.

b) Click **Reader** and click **Next**.

c) Click **Select Members**.

d) On the right side of the page, click the name of the app you registered and click **Select**.



e) Click **Review + Assign** and follow the prompts to complete the action.

Also see Create an Azure connector, on page 2639.

Create an Azure connector

Create a connector to send data from Azure to Security Cloud Control for use in policies.

This connector enables integration between Azure and your security management platform to retrieve dynamic object information for policy implementation.

Before you begin

Create an Azure user with at least the privileges discussed in Create an Azure user with minimal permissions for the dynamic attributes connector, on page 2637.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Azure. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Subscription ID	(Required.) Enter your Azure subscription ID.
Tenant ID	(Required.) Enter your tenant ID.
Client ID	(Required.) Enter your client ID.
Client Secret	(Required.) Enter your client secret.

a) Click **Test** and make sure **Test connection succeeded** is displayed before you save the connector.

b) Click **Save**.

The Status column displays **Ok**.

Create an Azure service tags connector

Create a connector for Azure service tags to enable use in policies with automatically updated IP addresses from Microsoft.

This topic discusses how to create a connector for Azure service tags to the Security Cloud Control for use in policies. The IP addresses associated with these tags are updated every week by Microsoft.

For more information, see Virtual network service tags on Microsoft TechNet.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter the required information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Azure. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Subscription ID	(Required.) Enter your Azure subscription ID.
Tenant ID	(Required.) Enter your tenant ID.
Client ID	(Required.) Enter your client ID.
Client Secret	(Required.) Enter your client secret.

a) Click **Test** and make sure **Test connection succeeded** is displayed before you save the connector.

b) Click **Save**.

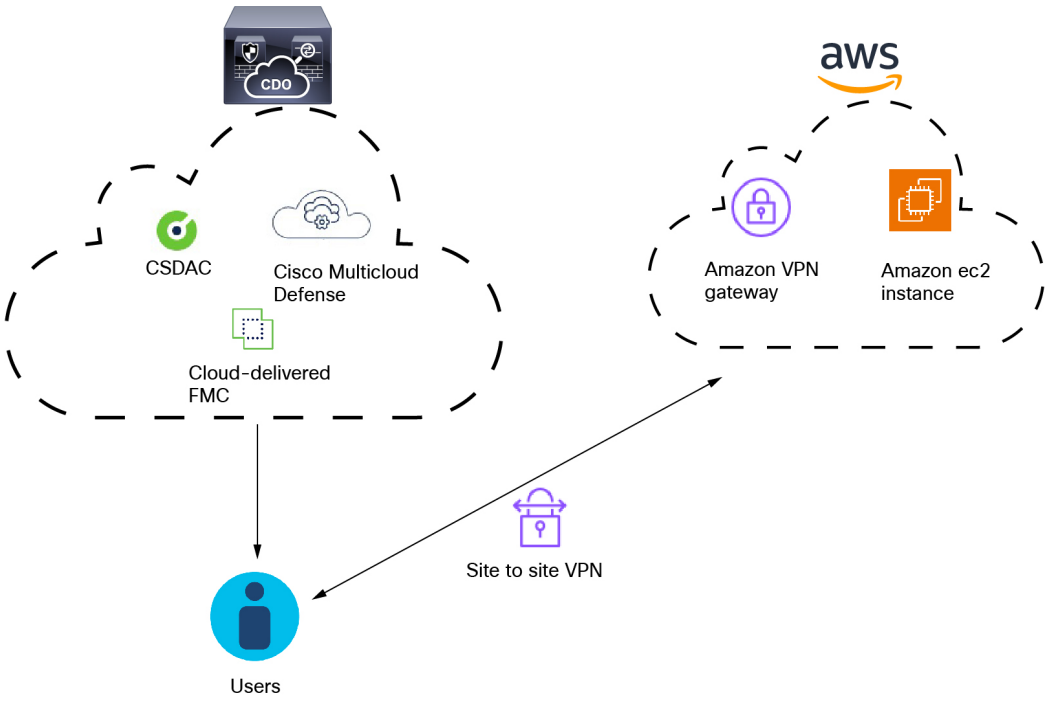
The Status column displays **Ok**.

Create a Multicloud Defense connector

This topic discusses how to create a connector for Cisco Multicloud Defense. The connector sends dynamic application address objects to the configured Cloud-Delivered Firewall Management Center.

For more information, see the Address Objects chapter in the *Cisco Multicloud Defense User Guide* and address object API documentation.

The following figure shows how the Cisco Multicloud Defense connector works.



As the figure shows:

- Users logging in and out of AWS create activity monitored by Multicloud Defense.
- The dynamic attributes connector and Multicloud Defense, both included in Security Cloud Control, send IP addresses from that activity to the Cloud-Delivered Firewall Management Center.
- These IP addresses can then be used in access control rules by the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter a **Name** and optional **Description** to identify the connector.

Step 6 Enter a **Pull Interval**. (Default 30 seconds.) Interval at which objects are retrieved from the Multicloud Defense Connector.

Step 7 Click **Test** and make sure the test succeeds before you save the connector.

Step 8 Click **Save**.

Step 9 Make sure **Ok** is displayed in the Status column.

You must create a Cloud-Delivered Firewall Management Center adapter as discussed in Create an adapter, on page 2659.

Create a Cisco cyber vision connector

This task creates a connector to send data from Cisco Cyber Vision to the firewall, enabling dynamic attribute integration between these systems.

This task discusses how to send data from Cisco Cyber Vision to the Security Cloud Control.

Before you begin

Cisco Cyber Vision must be reachable from the machine on which the dynamic attributes connector is running. You must know its IP address, port, and API key.

To find the API key in the Cyber Vision management console, click **Admin** > **API** > **Token**, then click **Show** to display the token and  to copy the token to the clipboard.

Follow these steps to create a Cisco Cyber Vision connector:

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon (), then click the name of the connector.
- Edit a connector: click Edit icon ().
- Delete a connector: click Delete icon ().

Step 5 Enter the required information.

Value	Description
Name	(Required) Enter a name to uniquely identify this connector.
Description	Optional description.

Value	Description
Cyber Vision Prefix	Enter an alphanumeric string to identify dynamic objects from this Cyber Vision's IP address when objects are sent to Security Cloud Control. If you have one Cyber Vision IP address, you can enter any value such as **1**.
Pull Interval	(Default 60 seconds) Interval at which data mappings are retrieved from Cyber Vision. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Host	(Required) Enter the Cyber Vision fully qualified host name or IP address.
Port	(Required) Enter the Cyber Vision listen port.
Version 4	(Required) Toggle this button to enable or disable Version 4. When you enable Version 4, the connector fetches hosts and vulnerabilities from the Cyber Vision platform. The Network Map API on the uses this data. When you disable Version 4, the connector fetches groups from the Cyber Vision platform and adds them as dynamic objects.
Token	(Required) Enter the API token.
Score Type	Choices are CVSS and CSRS. **CVSS**: Common vulnerability scoring system (CVSS) is an industry-standard system that retrieves values from the national vulnerability database to describe risk associated with vulnerabilities. CVSS scores power a vulnerability's severity and risk value. **CSRS**: The Cisco Security Risk Score (CSRS) CSRS evaluates vulnerabilities beyond technical severity, focusing on how attackers might exploit them. Scores range from 0 to 100 and are based on factors such as existing vulnerabilities, threat intelligence, and the effectiveness of security controls. This score helps prioritize critical vulnerabilities and allocate resources effectively.
Severity	Click the minimum vulnerability severity level for the dynamic attributes connector to send IP addresses to the Secure Firewall Management Center. (For example, if you click high, IP addresses of hosts with either high or severe vulnerabilities are sent.) Choices: • Critical • High • Medium • Low
Dynamic Object Name	Enter a name to identify the dynamic object created by this connector.

Value	Description
Certificate	You have these options: • Paste the certificate authority (CA) chain as discussed in Manually get a certificate authority (CA) chain, on page 2646. • Click **Get Certificate** > **Fetch** to automatically fetch the certificate. Ensure that you have specified a valid host and IP address that exactly matches the server certificate. • Click **Get Certificate** > **Browse from file** to upload a certificate chain you downloaded previously. Ensure that the host or IP address of the certificate you upload exactly matches the Common Name of the CA certificate you enter in host or IP address field.

a) Click **Test** and make sure the test succeeds before you save the connector.
b) Click **Save**.

Status column displays **Ok**.

Create an adapter, on page 2659

Create a generic text connector

Use this task to create a connector that retrieves IP addresses from text files and sends them to the firewall for use in access control rules.

This task discusses how to create an ad hoc list of IP addresses you maintain manually and retrieve at an interval you select (30 seconds by default). You can update the list of addresses anytime you want.

For example, you might have a list of IP addresses for an "allow list" in access control rules and another list of IP addresses for a "block list" in access control rules.

Before you begin

Create text files with IP addresses and put it on a web server that is accessible from the Security Cloud Control. IP addresses can include CIDR notation. The text file must have only one IP address per line.

You can specify up to 10,000 IP addresses per text file.

Follow these steps to create a generic text connector:

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter the following information:

Item	Description
Name	Enter a name to identify the connector.
Description	(Optional.) Enter a description
Pull Interval	Change the frequency, in seconds, at which the dynamic attributes connector retrieves IP addresses from the text file. The default is 30 seconds. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
URLs	Enter a URL from which to retrieve IP addresses.
Add another URL	(Optional.) Click the link to add more URLs to an existing list.
Certificate	(Optional.) If a certificate chain is required for a secure connection to the web server, you have these options: • Click **Get Certificate** > **Fetch** to automatically fetch the certificate or, if that is not possible, get the certificate manually as discussed in Manually get a certificate authority (CA) chain, on page 2646. • Click **Get Certificate** > **Browse from file** to upload a certificate chain you downloaded previously.

a) Click **Test** and make sure the test succeeds before you save the connector.
b) Click **Save**.
c) Status column displays **Ok**

Manually get a certificate authority (CA) chain

If the certificate authority chain is not fetched automatically, use one of these browser-specific procedures to get a certificate chain to securely connect to vCenter or Cloud-Delivered Firewall Management Center .

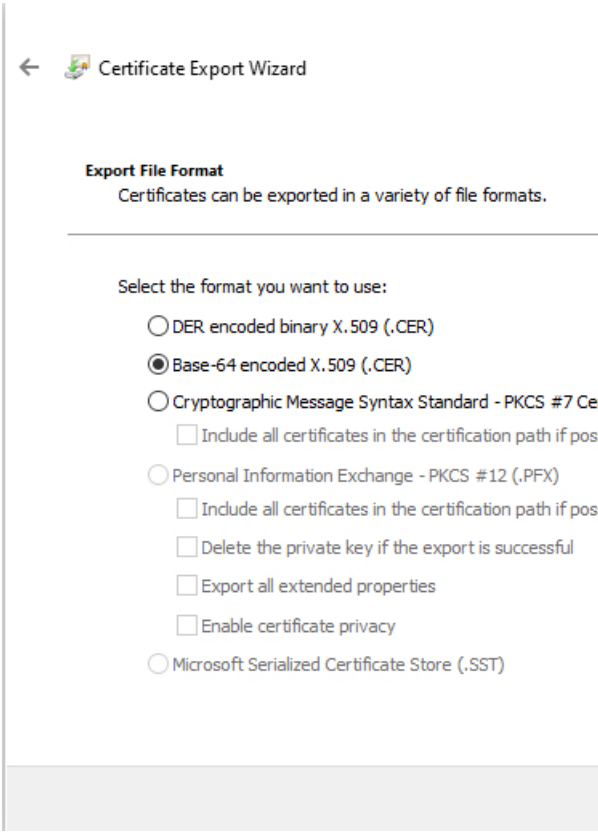
The *certificate chain* is the root certificate and all subordinate certificates.

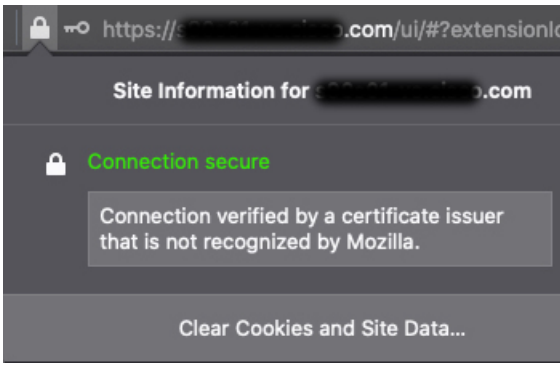
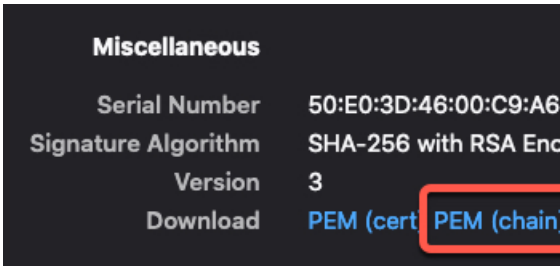
Procedure

	Command or Action	Purpose

Step 1	Get a Certificate Chain—Mac (Chrome and Firefox)	**a.** Open a Terminal window.
		b. Enter this command:
		`security verify-cert -P` *`url[:port]`*
		where *url* is the URL (including scheme) to vCenter or Cloud-Delivered Firewall Management Center . For example:
		`security verify-cert -P https://myvcenter.example.com`
		If you access vCenter or Cloud-Delivered Firewall Management Center using NAT or PAT, you can add a port as follows:
		`security verify-cert -P https://myvcenter.example.com:12345`
		c. Save the entire certificate chain to a plaintext file.
		• *Include* all `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` delimiters.
		• *Exclude* any extraneous text (for example, the name of the certificate and any text contained in angle brackets (< and >) as well as the angle brackets themselves.
		d. Repeat these tasks for vCenter Cloud-Delivered Firewall Management Center .

Step 2	Get a Certificate Chain—Windows Chrome	**a.** Log in to vCenter or Cloud-Delivered Firewall Management Center using Chrome. **b.** In the browser address bar, click the lock to the left of the host name. **c.** Click **Certificate**. **d.** Click the **Certification Path** tab. **e.** Click the top (that is, first) certificate in the chain. **f.** Click **View Certificate**.

	Command or Action	Purpose
		g. Click the **Details** tab. **h.** Click **Copy to File**. **i.** Follow the prompts to create a CER-formatted certificate file that includes the entire certificate chain. When you're prompted to choose an export file format, click **Base 64-Encoded X.509 (.CER)**.  Follow the prompts to complete the export. Open the certificate in a text editor. Repeat the process for all certificates in the chain. You must paste each certificate in the text editor in order, first to last. Repeat these tasks for vCenter or Cloud-Delivered Firewall Management Center .

	Command or Action	Purpose
Step 3	Get a Certificate Chain—Windows Firefox	**a.** Log in to vCenter or Cloud-Delivered Firewall Management Center using Firefox. **b.** Click the lock to the left of the host name. **c.** Click the right arrow (**Show connection details**).  **d.** Click **More Information**. **e.** Click **View Certificate**. **f.** If the resulting dialog box has tab pages, click the tab page corresponding to the top-level CA. **g.** Scroll to the Miscellaneous section. **h.** Click **PEM (chain)** in the Download row.  **i.** Save the file. Repeat these tasks for vCenter or Cloud-Delivered Firewall Management Center .

Create a GitHub connector

Create a GitHub connector that sends data to the Security Cloud Control for use in policies. The IP addresses associated with these tags are maintained by GitHub.

For more information, see About GitHub's IP addresses.

Note Do not change the URL because doing so will fail to retrieve any IP addresses.

Before you begin

You do not have to create a dynamic attributes filters.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter a **Name** and an optional description.

a) (Optional) In the **Pull Interval** field, change the frequency, in seconds, at which the dynamic attributes connector retrieves IP addresses from GitHub. The default is 21,600 seconds (6 hours).

b) Click **Test** and make sure the test succeeds before you save the connector.

c) Click **Save**.

The Status column displays **Ok**.

Google cloud connector user permissions and imported data

A Google Cloud connector imports dynamic attributes from Google Cloud to Security Cloud Control for use in policies.

Dynamic attributes imported

The connector imports these dynamic attributes from Google Cloud:

- *Labels*: Key-value pairs you can use to organize your Google Cloud resources. For more information, see Creating and Managing Labels in the Google Cloud documentation.
- *Network tags*: Key-value pairs associated with an organization, folder, or project. For more information, see Creating and Managing Tags in the Google Cloud documentation.
- *IP addresses*: Virtual machine addresses in Google Cloud.

The dynamic attributes connector requires a user at minimum with the **Basic** > **Viewer** permission to be able to import dynamic attributes.

Create a Google Cloud user with minimal permissions for the dynamic attributes connector

Create a service account with minimum permissions to send dynamic attributes to Security Cloud Control.

This task sets up a service account that provides the minimum permissions necessary for the dynamic attributes connector. For a list of these attributes, see Google cloud connector user permissions and imported data, on page 2650.

Before you begin

You must already have set up your Google Cloud account. For more information about doing that, see Setting Up Your Environment in the Google Cloud documentation.

Procedure

Step 1 Log in to your Google Cloud account as a user with the owner role.

Step 2 Click **IAM & Admin** > **Service Accounts** > **Create Service Account**.

Step 3 Enter the following information:

- **Service account name**: A name to identify this account; for example, **CSDAC**.
- **Service account ID**: Should be populated with a unique value after you enter the service account name.
- **Service account description**: Enter an optional description.

For more information about service accounts, see Understanding Service Accounts in the Google Cloud documentation.

Click **Create and Continue**.

Step 4 Follow the prompts on your screen until the Grant users access to this service account section is displayed.

Grant the user the **Basic** > **Viewer** role.

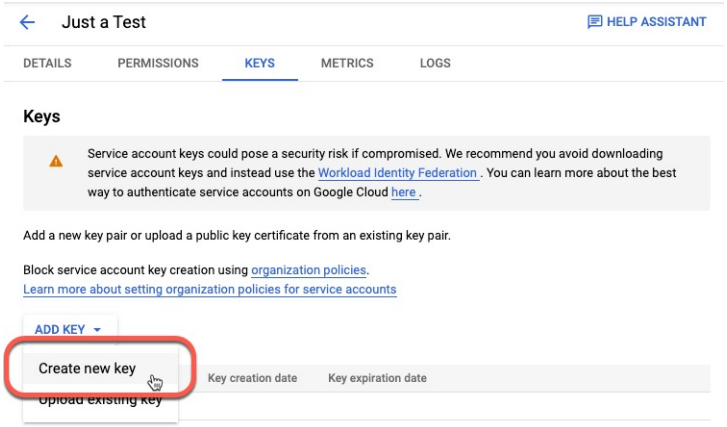
Click **Done**.

A list of service accounts is displayed.

Step 5 Click **More** (⋮) at the end of the row of the service account you created.

Step 6 Click **Manage Keys**.

Step 7 Click **Add Key** > **Create New Key**.



Step 8 Click **JSON**.

Step 9 Click **Create**.

The JSON key is downloaded to your computer.

Keep the key handy when you configure the GCP connector.

See Create a google cloud connector, on page 2652.

Create a google cloud connector

Create a Google Cloud connector to enable data transmission from your Google Cloud environment to the firewall for dynamic security attribute monitoring and management.

The Google Cloud connector allows the firewall to receive real-time data from your Google Cloud infrastructure, enabling dynamic security policies based on cloud resource attributes.

Before you begin

Have your Google Cloud JSON-formatted service account data ready; it's required to set up the connector.

Follow these steps to create a Google Cloud connector:

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from AWS. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
GCP region	(Required.) Enter the GCP region in which your Google Cloud is located. For more information, see Regions and Zones in the Google Cloud documentation.
Service account	Paste the JSON code for your Google Cloud service account.

a) Click **Test** and make sure the test succeeds before you save the connector.
b) Click **Save**.

The Status column displays **Ok**.

Create an office 365 connector

Create a connector for Office 365 tags to send data to the Security Cloud Control for use in policies.

The IP addresses associated with these tags are updated every week by Microsoft. You do not have to create a dynamic attributes filter to use the data.

For more information, refer to Office 365 URLs and IP address ranges on docs.microsoft.com.

Procedure

Step 1 Log in to Security Cloud Control.

Click **Firewall**.

Step 2 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon (+), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Azure. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Base API URL	(Required.) Enter the URL from which to retrieve Office 365 information, if it's different from the default. For more information, see Office 365 IP Address and URL web service on the Microsoft documentation site.
Instance name	(Required.) From the list, click an instance name. For more information, see Office 365 IP Address and URL web service on the Microsoft documentation site.
Disable optional IPs	(Required.) Enter **true** or **false**.

Click **Test** and make sure the test succeeds before you save the connector.

Step 5 Click **Save**.

The Status column displays **Ok**.

Create an adapter, on page 2659

Tenable connector

These topics provide information about creating and using the Tenable connector for vulnerability management.

About the Tenable connector

Tenable Vulnerability Management is a platform that helps organizations understand, report, and manage known vulnerabilities. When used with the Cloud-Delivered Firewall Management Center, the Tenable connector:

- Creates a dynamic object with a list of IP addresses known to have Common Vulnerability Exposures (CVEs).
- Populates those IP addresses and associated vulnerabilities to the Cloud-Delivered Firewall Management Center network map as host entries.

  Network map entries provide the basis for intrusion policy recommendations.

The dynamic object created by this connector can be used in access control rules, or anywhere dynamic objects are supported. For example, you can use the dynamic object to block vulnerable devices from accessing highly sensitive resources.

Supported products

We support Tenable Vulnerability Management only. We *do not* support Tenable Security Center.

Get the Tenable API key and secret

To configure the Tenable connector, you must create an API key and secret as described in this topic.



Note API keys inherit the permissions of the user account that generates them. Create the API key and secret using an account with full visibility to all hosts and vulnerabilities. For more information, see Permissions.

Procedure

Step 1 Log in to Tenable as an administrator.

Step 2 Click your profile image in the top right of the page.

Step 3 Click **My Profile** > **API Keys**.

Step 4 Click **Generate**.
You are required to confirm the action.

Step 5 Under Custom API keys, click (Copy to clipboard) to copy both the access key and secret key to the clipboard.

Step 6 Save these values for later use.

What to do next

See Create a Tenable connector, on page 2655.

Create a Tenable connector

This task explains how to configure the Tenable connector in Cisco Security Cloud. After you configure the connector, you must also configure an adapter to receive the dynamic object.

Before you begin

- Configure Tenable Vulnerability Management : Consult Tenable documentation.
- Get required information: Get the Tenable API key and secret, on page 2655.

We support Tenable Vulnerability Management only. We *do not* support Tenable Security Center.

Required User Role:

- Super Admin

Procedure

Step 1 Log in to Cisco Security Cloud.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 21600 seconds or six hours.) Interval at which IP mappings are retrieved from Tenable. We recommend a minimum value of 3600 seconds (one hour) to avoid issues with Tenable rate limiting.
Integration Key	Enter the API key you got in Get the Tenable API key and secret, on page 2655.
Secret Key	Enter the secret key obtained in Get the Tenable API key and secret, on page 2655.
Dynamic Object Name	Enter a name to identify the dynamic object created by this connector.
Severity Score	Click the minimum vulnerability severity level for the dynamic attributes connector to send IP addresses to the Cloud-Delivered Firewall Management Center. (For example, if you click **high**, IP addresses of hosts with either high or severe vulnerabilities are sent.) Choices: • **Critical** • **High** • **Medium** • **Low**

Value	Description
Severity System	Choices: • **VPR**: (Vulnerability priority rating.) Proprietary Tenable vulnerability rating that dynamically scores threats. VPR values range from 0.1-10.0, with a higher value representing a higher likelihood of exploit: • VPR **critical** is 9.0 and greater • VPR **high** is 7.0 and greater • VPR **medium** is 4.0 and greater • VPR **low** is 0.1 and greater • **CVSSv3**: (Common vulnerability scoring system version 3.) Industry-standard system that retrieves values from the national vulnerability database to describe risk associated with vulnerabilities. CVSS scores power a vulnerability's severity and risk value. For more information, see CVSS vs. VPR.

Step 6 Click **Test**. Save the connector only after the test succeeds.

Step 7 Click **Save**.

Step 8 Make sure **Ok** is displayed in the Status column.

See Create an adapter, on page 2659.

About Tenable dynamic objects in IDS, IPS, and access control policies

You can use IPS, IDS, and access control policies and rules to monitor or block traffic to and from servers with vulnerabilities identified by the Tenable connector:

1. To monitor traffic and inform you about vulnerabilities *without* blocking the traffic, create an intrusion detection system (IDS) policy with recommendations.
2. To monitor traffic, inform you about vulnerabilities, and block matching traffic, create an intrusion prevention system (IPS) policy with recommendations.
3. Create a new access control policy or add rules to an existing policy. Associate your IDS or IPS policy with an access control rule.

For information about intrusion policies for access control, refer to the Snort 3 configuration guide.

Create a Webex connector

This task creates a Webex connector that sends IP addresses to the management system for use in policies. The IP addresses associated with these tags are maintained by Webex.

This section discusses how to create a Webex connector that sends data to the Security Cloud Control for use in policies. The IP addresses associated with these tags are maintained by Webex. You do not have to create a dynamic attributes filters.

For more information, see Port Reference for Webex Calling.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Webex. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Provider Reserved IPs	(Required.) (Required.) Slide to enabled to retrieve any reserved IP addresses.

a) Click **Test** and make sure the test succeeds before you save the connector.
b) Click **Save**.
c) The Status column displays **Ok**.

Create a Zoom connector

Create a Zoom connector that sends data to the Security Cloud Control for use in policies.

The IP addresses associated with these tags are maintained by Zoom. You do not have to create a dynamic attributes filters.

For more information, see Zoom network firewall or proxy server settings.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit a connector: click Edit icon (Edit).
- Delete a connector: click Delete icon (Delete).

Step 5 Enter this information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Zoom. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Provider Reserved IPs	(Required.) Slide to enabled to retrieve any reserved IP addresses.

a) Click **Test** and make sure the test succeeds before you save the connector.

b) Click **Save**.

The Status column displays **Ok**.

Create an adapter

An *adapter* is a secure connection to Cloud-Delivered Firewall Management Center or an On-Prem Firewall Management Center to which you push network information from cloud objects for use in access control policies.

You can create the following adapters:

- *On-Prem Firewall Management Center* for an on-premises Cloud-Delivered Firewall Management Center.

- *Cloud-Delivered Firewall Management Center* for devices managed by Cisco Security Cloud.

Note You must have a **Super Admin** user role to create the first adapter. To view or modify existing adapters, you must have an Admin or Super Admin user role.

How to create an On-Premises Firewall Management Center adapter

This topic discusses how to create an adapter to push dynamic objects from the dynamic attributes connector to Cisco Security Cloud.

Before you begin

Onboard the firewall manager to Cisco Security Cloud as discussed in *Onboard a Management Center* in the *Managing Security and Network Devices with Security Cloud Control* online help.

Note Currently, you can use the Tenable connector with the On-Premises Firewall Management Center only to create the dynamic object. The On-Premises Firewall Management Center does *not* receive network map entries from the Tenable connector, so you cannot use it with intrusion policy recommendations.

Required User Role: Super Admin

Procedure

Step 1 Log in to Cisco Security Cloud as a user with the Super Admin role.

Step 2 Click **Administration** > **Dynamic Attributes Connector** > **Adapters**.

Step 3 To add an adapter, click Add icon ( ) > On-Premises Firewall Management Center.

Step 4 To edit or delete an adapter, click Edit icon (Edit), or Delete icon (Delete).

Step 5 Add or edit this information.

Value	Description
Name	(Required.) Enter a unique name to identify this adapter.
Description	(Optional) Enter a description of the adapter.
Primary Device	Click the IP address of a management center associated with your tenant.
Secondary Device	(Optional) If you have a secondary On-Premises Firewall Management Center, click its name.

Step 6 Click **OK**.

How to create a Cloud-Delivered Firewall Management Center adapter

This topic discusses how to create an adapter to push dynamic objects from the dynamic attributes connector to Cisco Security Cloud.

Before you begin

Required User Role:

- Super Admin

Procedure

Step 1 Log in to Cisco Security Cloud as a user with the Super Admin role.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Adapters**.

Step 4 To add an adapter, click Add icon ( ) > Cloud-Delivered Firewall Management Center.

Step 5 To edit or delete an adapter, click Edit icon (Edit), or Delete icon (Delete).

Step 6 Edit the following information.

Value	Description
Name	(Required.) Enter a unique name to identify this adapter.
Description	Optional description of the adapter.
Cloud FMC URL	From the list, click the URL for your Cloud-Delivered Firewall Management Center.

Step 7 Click **Test** and make sure the test succeeds before you save the adapter.

Step 8 Click **Save**.

Create dynamic attributes filters

Create dynamic attributes filters to expose them as dynamic objects in Security Cloud Control for use in access control policies.

Dynamic attributes filters that you define using the Dynamic Attributes Connector are exposed in the Security Cloud Control as dynamic objects that can be used in access control policies. For example, restrict access to an AWS server for the Finance Department to only members of the Finance group defined in Microsoft Active Directory.

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

For more information about access control rules, see Create access control rules or DNS rules using dynamic attributes filters, on page 2456.

Before you begin

Connectors, on page 2631

Procedure

Step 1 Log in to Security Cloud Control.

a) Click **Firewall**.

b) Click **Administration** > **Dynamic Attributes Connector** > **Dynamic Attributes Filters**.

Step 2 Do any of the following:

- Add a new filter: click **Add** ( ).
- Edit or delete a filter: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 3 Enter the following information.

Item	Description
Name	Unique name to identify the dynamic filter (as a dynamic object) in a policy and in the Security Cloud Control Object Manager (**External Attributes** > **Dynamic Object**).
Connector	From the list, click the name of a connector to use.
Query	Click Add .

Step 4 To add or edit a query, enter the following information.

Item	Description
Key	Click a key from the list. Keys are fetched from the connector.
Operation	Click one of the following: • **Equals** to exactly match the key to the value. • **Contains** to match the key to the value if any part of the value matches.

Item	Description
Values	Click either **Any** or **All** and click one or more values from the list. Click **Add another value** to add values to your query.

Step 5 Click **Show Preview** to display a list of networks or IP addresses returned by your query.

Step 6 When you're finished, click **Save**.

Step 7 (Optional.) Verify the dynamic object in the Security Cloud Control.

a) Log in to the Security Cloud Control.

b) Click **Policies** > **Firewall Threat Defense**.

c) Click **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.

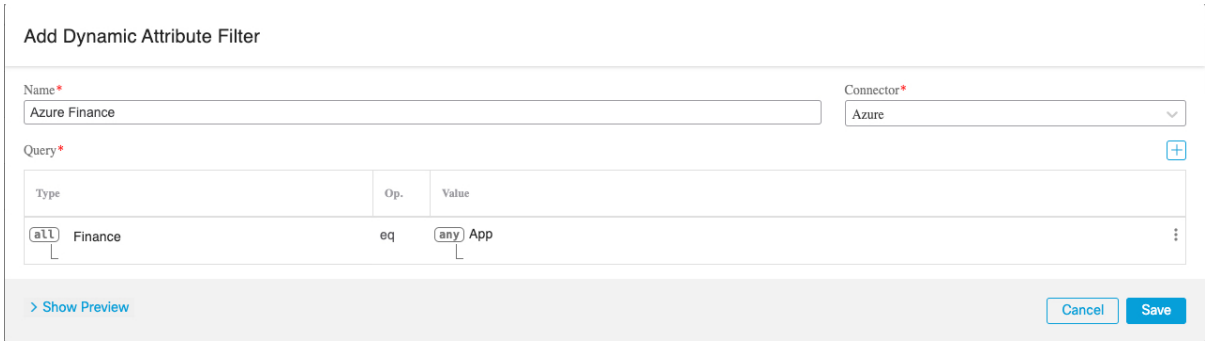
The dynamic attribute query you created should be displayed as a dynamic object.

Dynamic attribute filter examples

This topic provides some examples of setting up dynamic attribute filters.

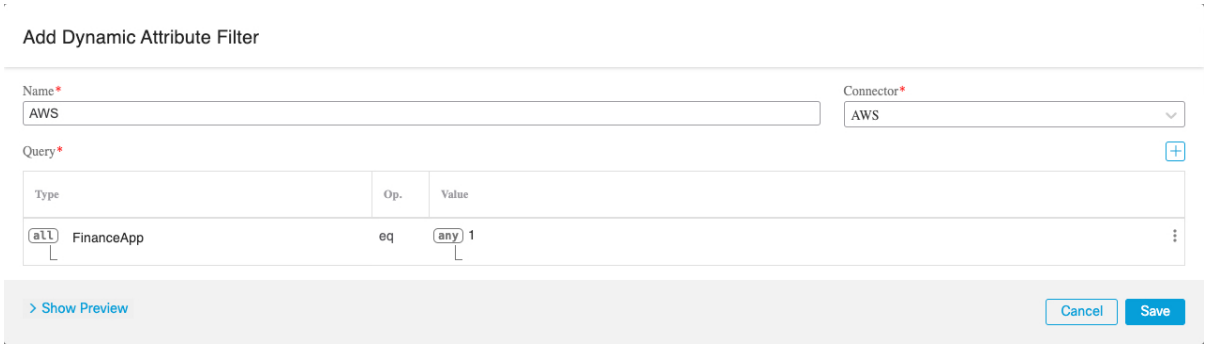
Example: azure

This example shows one criterion: a server tagged as a Finance app.



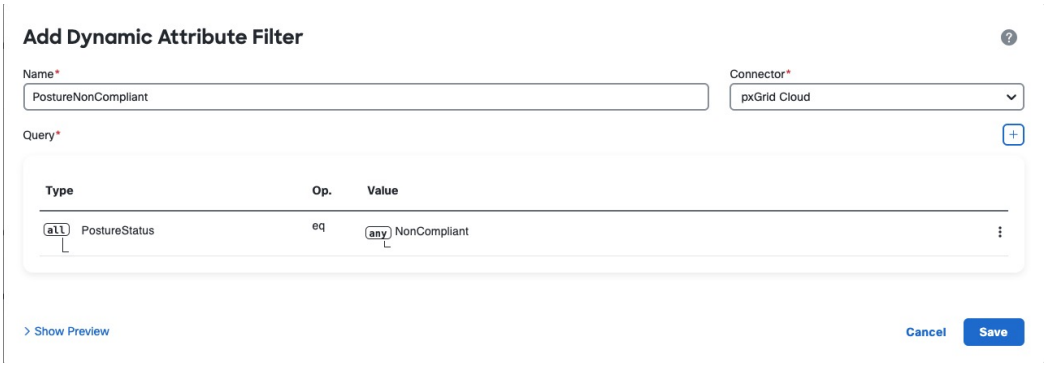
Example: AWS

This example shows one criterion: a FinanceApp with a value of 1.



Example: pxGrid cloud

This example shows one criterion: PostureStatus is NonCompliant.



Disable the dynamic attributes connector

Disable the Dynamic Attributes Connector when you no longer wish to collect dynamic objects from cloud sources.

If you no longer wish to collect dynamic objects from cloud sources, you can disable the Dynamic Attributes Connector in the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Log in to the Cloud-Delivered Firewall Management Center if you have not done so already.

Step 2 Click **Administration** > **Dynamic Attributes Connector**.

Step 3 Slide to **Disabled**.

Dynamic firewall

A dynamic firewall integrates user identity data with user trust data provided by Identity Intelligence to enhance your ability to detect identity-based exploits in your network.

Identity data integration

The dynamic firewall integrates user identity data from Microsoft AD and ISE with user trust data.

Dynamic Firewall

A dynamic firewall integrates identity sources with Identity Intelligence to provide user and device trust information.

Enhanced identity information

Previously, the Security Cloud Control collected information about users exclusively from the configured identity sources, such as Microsoft Active Directory, the passive identity agent, Cisco Identity Services Engine (Cisco ISE), and so on. This information generally included user name, group, and IP address.

The dynamic firewall enables you to pair user identity with intelligence and use that information in reporting and access control policies.

To use the dynamic firewall, you must:

- Have an Identity Intelligence tenant

  See Duo Identity Security with Cisco Identity Intelligence.

- Set up an identity source:
  - Cisco Identity Services Engine (Cisco ISE)
  - pxGrid Cloud

    pxGrid Cloud combines identity and posture in the same feed

    More information: What is pxGrid?

  In addition to providing authentication information, Cisco ISE and pxGrid Cloud can provide the following:

  - SGT Exchange Protocol (SXP) over TCP binding and directory session information if desired. For more information, see the *Cisco Identity Services Engine Administrator Guide*
  - Posture and mobile device management compliance. For more information, see Compliance.

- Set up an identity realm:
  - Create an LDAP realm or an Active Directory realm and realm directory, on page 2325
  - Create a Microsoft Azure AD (SAML) realm for passive authentication, on page 2298

The *identity source* provides authentication information (login, logout) as well as posture. The identity source can also provide SXP binding and session directory information if desired.

The *identity realm* provides user, group, and IP address information.

How to configure the dynamic firewall

This topic helps you understand the concepts and options to configure the dynamic firewall discussed in Dynamic Firewall, on page 2664.

Summary

The dynamic firewall integrates an identity source (such as Cisco ISE) with Cisco Identity Intelligence, which provides user trust information to the Cloud-Delivered Firewall Management Center.

1. Configure Cisco Identity Intelligence to collect user trust information.
2. Configure a supported Cloud-Delivered Firewall Management Center identity source.
3. Configure a supported identity realm.

4. Enable the dynamic attributes connector.
5. Configure the dynamic firewall.

Workflow

The following procedure provides a high-level overview of how to configure the dynamic firewall.

1. As a Duo user with the Owner role, provision a Cisco Identity Intelligence tenant.

   You can provision a tenant from Duo Advantage as discussed in *Provision Your Cisco Identity Intelligence Tenant*.

2. In Cisco Identity Intelligence, create an API integration and use the information to set up the dynamic firewall.

   We use Cisco Identity Intelligence to find user and device risk information in your network.

   For more information about Cisco Identity Intelligence, see How-to Guides.

   For more information about this task, see Get required information for Identity Intelligence, on page 2667.

3. (Microsoft Azure AD realm only.) In Identity Intelligence, create a Microsoft Entra ID integration.

   For more information, see Microsoft Entra ID (Azure AD) Data Integration.

4. Create an identity source. (If you already have an identity source, continue with the next step.)

   You can do this in any of the following ways:

   - The Configure Dynamic Firewall dialog box displays **Configure** links to start setting up your identity source.
   - Click **Integration** > **Other Integrations** > **Identity Sources**.

   For more information about creating identity sources, see:

   - Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source, on page 2390
   - Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436
   - Configuring a pxGrid Cloud identity source, on page 2468

5. Create an identity realm.

   We support the following realms:

   - Create an LDAP realm or an Active Directory realm and realm directory, on page 2325

     Only Microsoft AD is supported; LDAP realms are not supported.

   - Create a Microsoft Azure AD (SAML) realm for passive authentication, on page 2298

6. Create the dynamic firewall instance. (If you already have a dynamic firewall instance, continue with the next step.)

   Click **Administration** > **Dynamic Attributes Connector** and click **Configure Dynamic Firewall**.

   See Create a dynamic firewall instance, on page 2669.

7. Associate your identity source with Cisco Identity Intelligence.

See Associate an identity source with Identity Intelligence, on page 2670.

8. View system-defined filters.

   We create dynamic attributes filters for the following:

   - Untrusted device
   - Trusted device
   - Untrusted user
   - Questionable user

   You can edit or replace these dynamic attributes filters as discussed in Create dynamic attributes filters, on page 2677.

9. View system-defined access control rules.

   We create an access control policy named Dynamic Firewall Policy (or similar) with the following rules:

   - Block an untrusted user from any source network to any destination network.
   - Monitor a questionable user from any source network to any destination network.
   - Block an untrusted device from any source network to any destination network.

   You can edit or delete the access control policy and rules as discussed in View and edit the system-created access control policy, on page 2676.

Get required information for Identity Intelligence

Get the necessary credentials and configuration values from Identity Intelligence to enable integration with the dynamic firewall.

This task discusses how to create an API client, which provides all the get required information to set up Identity Intelligence in the dynamic firewall.

If you already have an API client and you know the values of all these parameters, you can skip this procedure and continue with Create a dynamic firewall instance, on page 2669:

- **Client ID**
- **API URL**
- **Token URL**
- **Client Secret**

Before you begin

Integrating with the dynamic firewall requires you to create an *API client integration* in Identity Intelligence.

Among the values you must know about your API client integration is the client secret, which is displayed when you create the API client only. For that reason you might need to create the API integration first.

For more information about creating an API cilent integration, see Public API.

Follow these steps to get required information for Identity Intelligence:

Procedure

Step 1 Log in to your Identity Intelligence tenant.

Step 2 Click (**Integrations**).

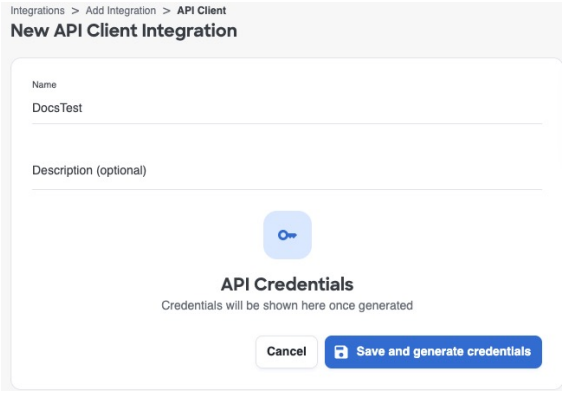
Step 3 Click **Add Integration**.

Step 4 On the next page, under API Clients, click **Add API Client**.

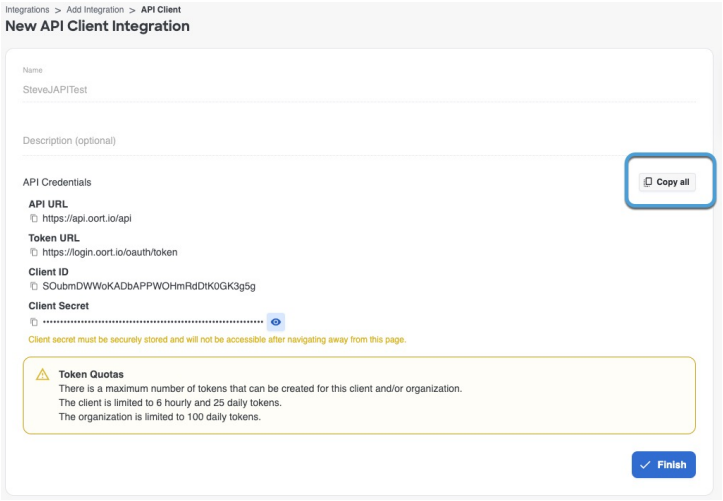
Step 5 Enter a **Name** and an optional **Description**.

Step 6 Click **Save and Generate Credentials**.

This figure shows an example.



Step 7 On the next page, click **Copy all** as this figure shows.



Step 8 Save the credentials for later use.

Step 9 Click **Finish**.

What to do next

See Create a dynamic firewall instance, on page 2669.

Create an identity source and realm for the dynamic firewall

Before you configure the dynamic firewall, you must configure a supported identity realm and identity source.

Configure an identity realm

These identity realms are supported:

- Create an LDAP realm or an Active Directory realm and realm directory, on page 2325

  Only Microsoft AD is supported; LDAP realms are not supported.

- Create a Microsoft Azure AD (SAML) realm for passive authentication, on page 2298

Configure an identity source

These identity sources are supported:

- On-premises Cisco ISE: Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source, on page 2390
- Single or multiple Cisco ISE clusters:
  - Configuring a pxGrid Cloud identity source, on page 2468
  - Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436

Create a dynamic firewall instance

Create a new instance of the dynamic firewall to establish an association between an identity source and Identity Intelligence.

This task discusses how to create a new instance of the dynamic firewall, which is an association between an identity source and Identity Intelligence.

Before you begin

Complete all these steps:

- Create an identity source:
  - Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source, on page 2390.
  - Create a pxGrid Cloud identity source, on page 2443.

Follow these steps to create a dynamic firewall instance:

Procedure

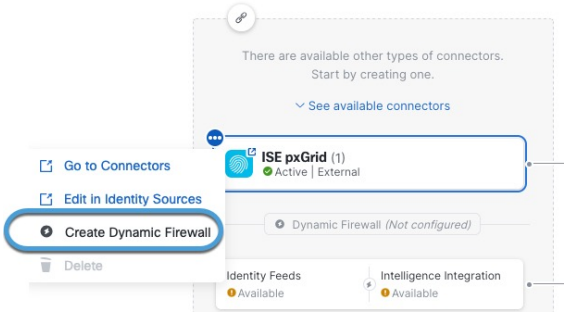
Step 1 If you haven't already done so, log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector**.

Step 4 Click next to the name of the identity source with which to add the dynamic firewall.

The following figure shows an example.



Note
If you do not see an identity source, create one before continuing:

- Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source, on page 2390
- Configuring a pxGrid Cloud identity source (Cisco ISE 3.3 or earlier), on page 2436
- Configuring a pxGrid Cloud identity source, on page 2468

Step 5 Click **Create Dynamic Firewall**.

Step 6 Continue with Associate an identity source with Identity Intelligence, on page 2670.

Associate an identity source with Identity Intelligence

Associate a supported identity source with Identity Intelligence to enable user and device trust ratings for dynamic firewall configuration.

This task discusses how you associate an identity source with Identity Intelligence, which provides user and device trust ratings to the Security Cloud Control.

For more information, see User Trust Level.

Before you begin

Before you begin, make sure you:

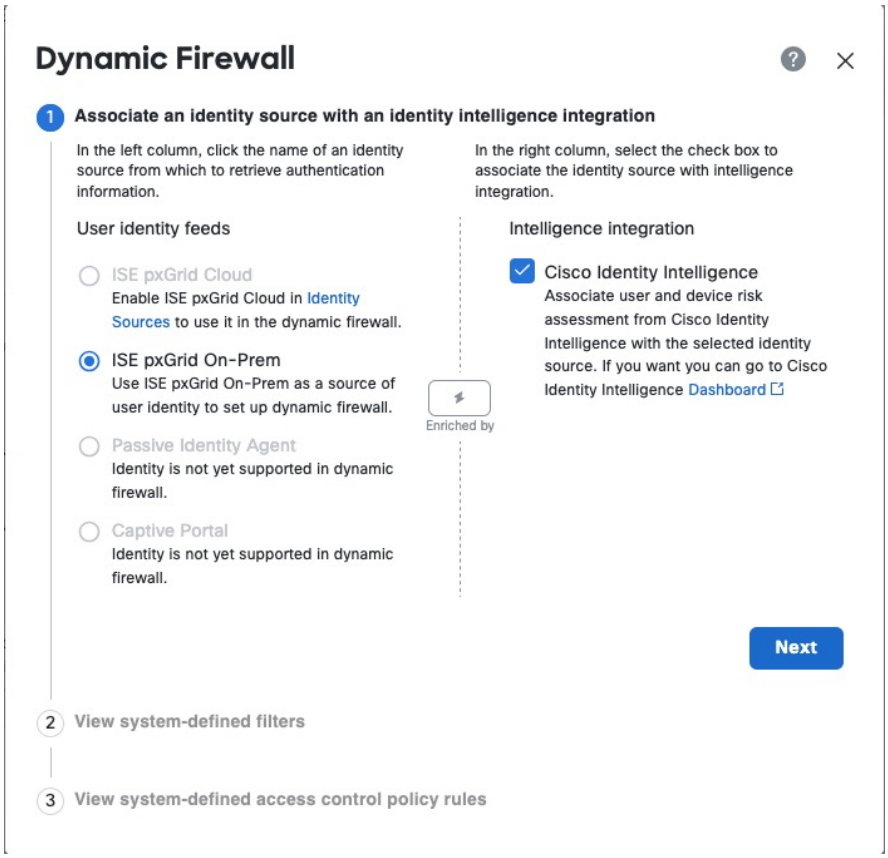
- Understand how the identity realm, identity source, and Identity Intelligence work together as discussed in Dynamic Firewall, on page 2664.
- Completed the tasks discussed in Create a dynamic firewall instance, on page 2669.

Procedure

Step 1 Start with Create a dynamic firewall instance, on page 2669.

Step 2 On the next page, from the left column, click your identity source. Then, in the right column, select the **Cisco Identity Intelligence** check box to add user intelligence, including user and device risk.

The following figure shows an example.



Step 3 Click **Next**.

Step 4 Continue with Configure Identity Intelligence, on page 2671.

Configure Identity Intelligence

This task enables integration between the dynamic firewall and Identity Intelligence.

This task discusses how you associate an identity source with Identity Intelligence, which provides user and device risk ratings to the Security Cloud Control.

Before you begin

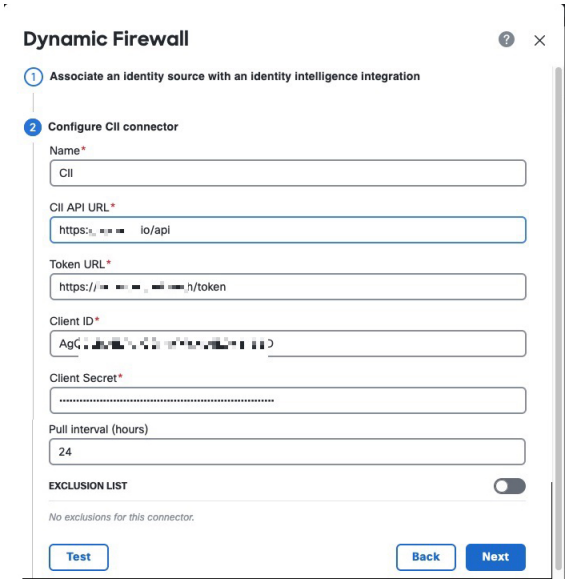
Complete the tasks discussed in Associate an identity source with Identity Intelligence, on page 2670.

Procedure

Step 1 Complete the tasks discussed in Associate an identity source with Identity Intelligence, on page 2670.

Step 2 If you selected the **Cisco Identity Intelligence** check box, enter the information you found for Identity Intelligence as described in Get required information for Identity Intelligence, on page 2667.

The following figure shows an example.



Step 3 (Optional.) For Identity Intelligence to consider a specific set of users as trusted, slide **Exclusion List** to **Slider enabled** ( ).

Enter one user name per line in **username@domain.com** format. Users in this list are considered trusted by Identity Intelligence.

Step 4 Click **Test**.

Only if the test succeeds, continue with the next step.

If any errors are displayed, check all of your Identity Intelligence values and try again.

Step 5 Click **Next**.

Step 6 Continue with View system-defined filters, on page 2672.

View system-defined filters

This task allows you to examine and understand the system-defined dynamic attributes filters that are automatically created during the configuration process.

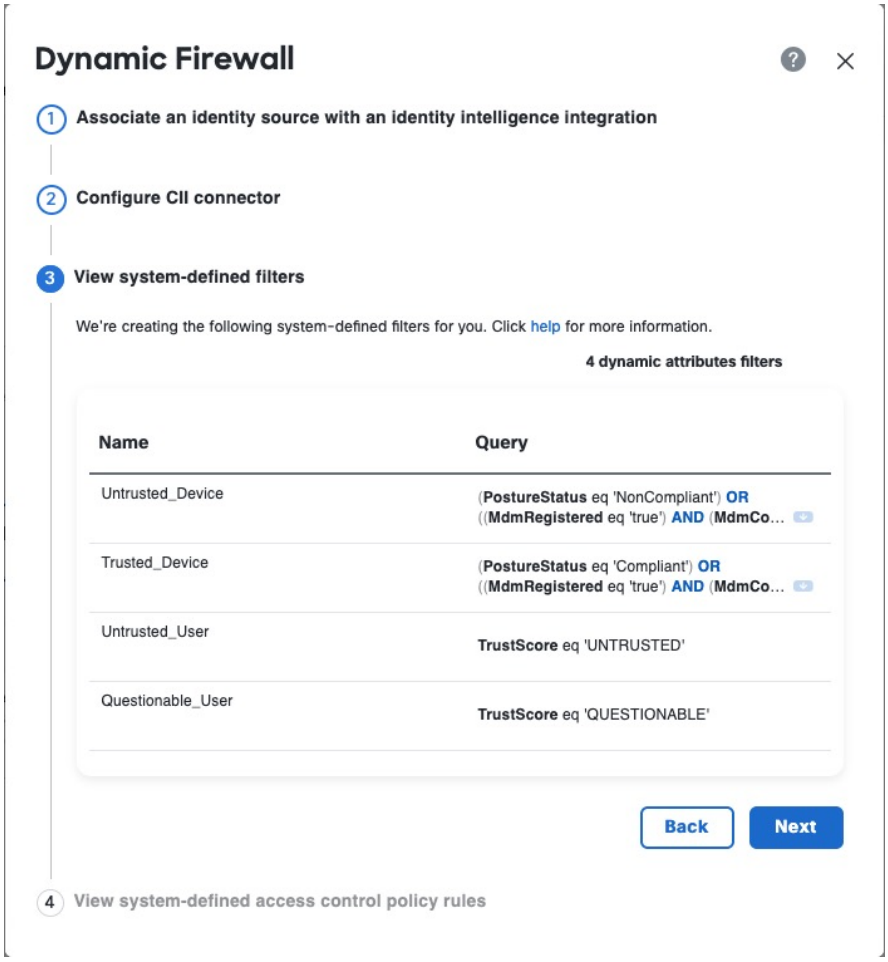
This task discusses how you associate an identity source with Cisco Identity Intelligence, which provides user and device risk ratings to the Security Cloud Control.

Before you begin

See Configure Identity Intelligence, on page 2671.

Procedure

Step 1 The system displays a set of system-defined dynamic attributes filters, as shown in the following figure.



Step 2 View the system-created filters. Click on any row to expand the filter so you can view the filter and see its details.

Step 3 Click **Next**.

Step 4 Continue with View system-defined access control rules, on page 2673.

View system-defined access control rules

This task allows you to review and decide whether to create system-defined access control rules for your dynamic firewall instance.

This task discusses access control rules created by the dynamic firewall.

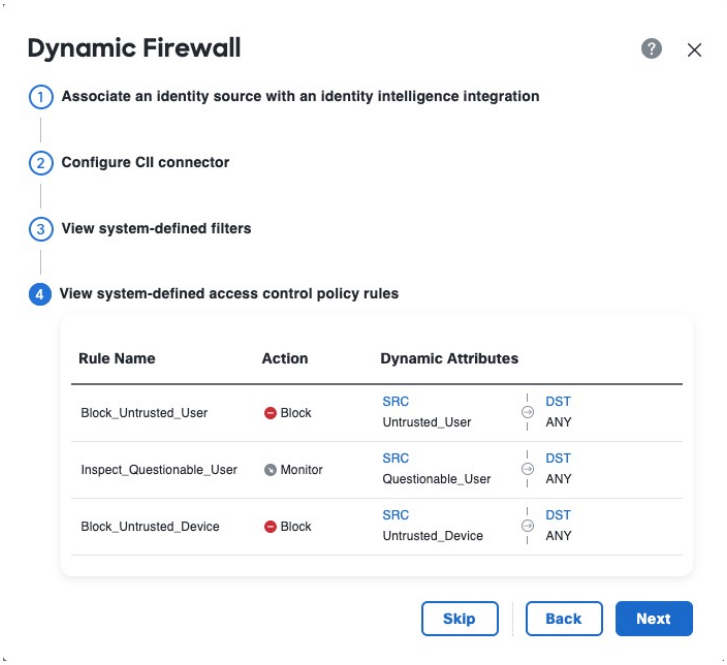
Before you begin

See View system-defined filters, on page 2672.

Procedure

Step 1 View the system-created access control rules.

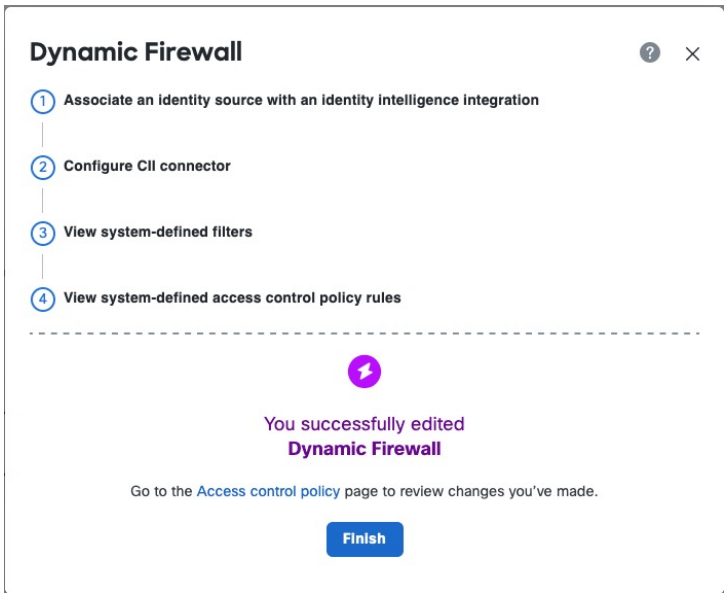
The following figure shows an example.



Step 2 Choose one of these options:

- Click **Skip** to skip creating these access control rules. You can create your own anytime.
- Click **Next** to create an access control policy named Dynamic Firewall Policy with the rules shown in the preceding figure.
- Click **Back** to return to system-created filters.

Step 3 After you click **Next**, if you created access control rules successfully, the following page is displayed:



Edit the user exclusion list

Specify users provided by Identity Intelligence as always being trusted.

You can instruct Identity Intelligence to treat specific users as trusted.

Before you begin

Configure the dynamic firewall as discussed in Create a dynamic firewall instance, on page 2669.

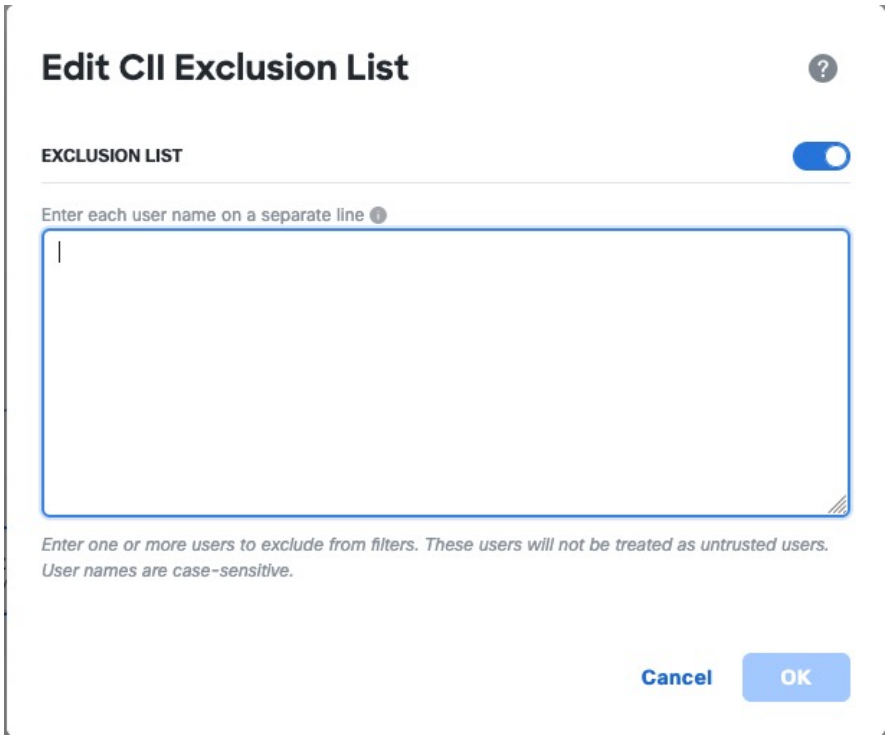
Procedure

Step 1 Click **Administration** > **Dynamic Attributes Connector**.

Step 2 Click next to the name of the identity source.

Step 3 Click **Edit CII Exclusion List**.

The following dialog box is displayed.



Step 4 In the provided field, enter one user name in `username@domain.com` format on a line, press Enter, and enter another user name.

Each user name is considered as trusted by Identity Intelligence.

View and edit the system-created access control policy

Customize the system-created access control policy by adding devices, changing rules, reordering rules, or deleting rules to meet your specific network security requirements.

This topic discusses how you can edit the system-created access control rules and policy. Initially, the policy isn't associated with any devices but if you want to use it you can add devices, change rules, reorder rules, or delete rules.

Before you begin

Complete the tasks described in View system-defined access control rules, on page 2673.

Follow these steps to view and edit the system-created access control policy:

Procedure

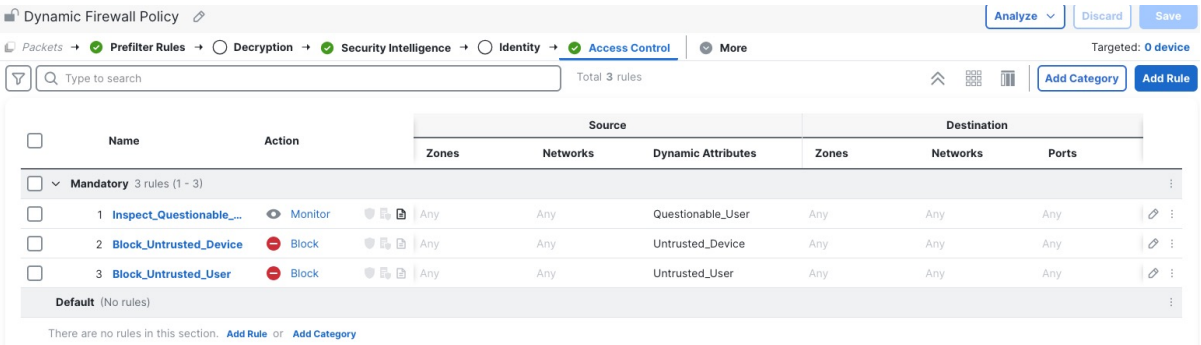
Step 1 If you haven't already done so, log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Policies** > **Firewall Threat Defense** > **Access Control heading** > **Access Control**.

Step 4 Click **Edit** ( ) next to the policy named Dynamic Firewall Policy (or similar).

The following figure shows a sample access control policy.



Note that in this access control policy, only the rule set to monitor questionable users logs anything. To adjust the logging settings, see Logging settings for access control policies, on page 2014.

Step 5 Do any of the following:

- Target the access control policy at devices: Assign access control policy to devices, on page 2014.
- Edit the policy, including adding logging: Manage access control policies, on page 2002.
- Edit access control rules: Access control rule management, on page 2062.
- Set advanced policy options: Configure advanced settings for the access control policy, on page 2688.
- Associate other policies with this access control policy: Associate prefilter, decryption, and identity policies with an access control, on page 2016.

Create dynamic attributes filters

Create dynamic attributes filters that can be used as dynamic objects in access control policies to restrict access based on specific criteria such as group membership.

Dynamic attributes filters that you define using the Dynamic Attributes Connector are exposed in the Security Cloud Control as dynamic objects that can be used in access control policies. For example, you could restrict access to an AWS server for the Finance Department to only members of the Finance group defined in Microsoft Active Directory.

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

For more information about access control rules, see Create access control rules or DNS rules using dynamic attributes filters, on page 2456.

Procedure

Step 1 Log in to Security Cloud Control.

Step 2 Click **Firewall**.

Step 3 Click **Administration** > **Dynamic Attributes Connector** > **Connectors**.

Step 4 Do any of the following:

- Add a new filter: click **Add** ( ).
- Edit or delete a filter: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 5 Enter the following information.

- **Name**: Unique name to identify the dynamic filter (as a dynamic object) in access control policy and in the Security Cloud Control Object Manager (**External Attributes** > **Dynamic Object**).
- **Connector**: From the list, click the name of a connector to use.

Step 6 To add a query, enter the following information.

- **Key**: Click a key from the list. Keys are fetched from the connector.
- **Operation**: Click one of the following:
  - **Equals** to exactly match the key to the value.
  - **Contains** to match the key to the value if any part of the value matches.
- **Values**: Click either **Any** or **All** and click one or more values from the list. Click **Add another value** to add values to your query.

Step 7 Click **Show Preview** to display a list of networks or IP addresses returned by your query.

When you're finished, click **Save**.

What to do next

Verify the dynamic object in the Security Cloud Control.

1. Log in to the Security Cloud Control.
2. Click **Policies** > **Firewall Threat Defense**.
3. Click **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.

The dynamic attribute query you created should be displayed as a dynamic object.

Use Dynamic Objects in Access Control Policies

The dynamic attributes connector enables you to configure dynamic attributes filters, seen in the Security Cloud Controlas dynamic objects, in access control rules.

Dynamic objects in access control rules

A *dynamic object* is automatically pushed from the dynamic attributes connector to the Secure Firewall Manager after you create connectors and save a dynamic attributes filter on the connector.

You can use these dynamic objects on the access control rule's **Dynamic Attributes** tab page. You can add dynamic objects as source or destination attributes.

Dynamic object usage examples

In an access control block rule, you can add a Finance dynamic object as a destination attribute to block access to Finance servers by whatever objects match the other criteria in the rule.

Dynamic attributes filter limitations

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

Create dynamic attributes filters

Create dynamic attributes filters to expose them as dynamic objects in Security Cloud Control for use in access control policies.

Dynamic attributes filters that you define using the Dynamic Attributes Connector are exposed in the Security Cloud Control as dynamic objects that can be used in access control policies. For example, restrict access to an AWS server for the Finance Department to only members of the Finance group defined in Microsoft Active Directory.

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

For more information about access control rules, see Create access control rules or DNS rules using dynamic attributes filters, on page 2456.

Before you begin

Connectors, on page 2631

Procedure

Step 1 Log in to Security Cloud Control.

a) Click **Firewall**.

b) Click **Administration** > **Dynamic Attributes Connector** > **Dynamic Attributes Filters**.

Step 2 Do any of the following:

- Add a new filter: click **Add** ( ).
- Edit or delete a filter: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 3 Enter the following information.

Item	Description
Name	Unique name to identify the dynamic filter (as a dynamic object) in a policy and in the Security Cloud Control Object Manager (**External Attributes** > **Dynamic Object**).
Connector	From the list, click the name of a connector to use.
Query	Click Add .

Step 4 To add or edit a query, enter the following information.

Item	Description
Key	Click a key from the list. Keys are fetched from the connector.
Operation	Click one of the following: • **Equals** to exactly match the key to the value. • **Contains** to match the key to the value if any part of the value matches.
Values	Click either **Any** or **All** and click one or more values from the list. Click **Add another value** to add values to your query.

Step 5 Click **Show Preview** to display a list of networks or IP addresses returned by your query.

Step 6 When you're finished, click **Save**.

Step 7 (Optional.) Verify the dynamic object in the Security Cloud Control.

a) Log in to the Security Cloud Control.

b) Click **Policies** > **Firewall Threat Defense**.

c) Click **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.

The dynamic attribute query you created should be displayed as a dynamic object.

Dynamic attributes rule conditions

Dynamic attributes rule conditions allow you to specify access control rules based on real-time properties such as network objects, device types, locations, and security tags. These conditions determine how access control rules match traffic based on various attributes, optimizing policy flexibility and accuracy.

Types of dynamic attributes

Dynamic attributes include:

- (source or destination) Dynamic objects (such as from the dynamic attributes connector)

  The dynamic attributes connector enables you to collect data (such as networks and IP addresses) from cloud providers and send it to the Cloud-Delivered Firewall Management Center so they can be used in access control rules.

  For more information about the dynamic attributes connector, see About the Dynamic Attributes Connector , on page 2621.

- (source only) SGT objects contain tags either manually defined or defined in ISE. For more information, see Source and destination security group tag (SGT) matching, on page 2374 and Security group tags, on page 2972.

- (source only) Location IP objects, defined by Cisco ISE

- (source only) Device type objects, defined by Cisco ISE (also referred to as endpoint profile objects)

Combining objects in rules

Dynamic attributes can be used as source criteria and destination criteria in access control rules. Use these guidelines:

- Objects of different types are ANDd together

- Objects of a similar type are ORd together

For example, if you choose source destination criteria SGT 1, SGT 2, and device type 1, the rule matches if device type 1 is detected on either SGT 1 or SGT 2.

If you select both a security group tag, and a dynamic object that lists IP addresses, the rule matches if traffic with the tag originating from, or destined to, one of those IP addresses.

Create access control rules or DNS rules using dynamic attributes filters

Create access control rules using dynamic objects that are based on dynamic attributes filters you have previously configured.

This topic discusses how to create access control rules using dynamic objects (these dynamic objects are named after the dynamic attributes filters you created previously).

To add dynamic attributes filters to DNS policies, see Create basic DNS policies, on page 2561.

To add dynamic attributes filters to DNS policies, see Creating Basic DNS Policies.

Before you begin

Create dynamic attributes filters as discussed in Create dynamic attributes filters, on page 2455.

Note You cannot create dynamic attributes filters for Azure Service Tags, Cisco Multicloud Defense, Generic Text, GitHub, Outlook 365, Tenable, Webex, and Zoom). These types of cloud objects provide their own IP addresses.

Follow these steps to create access control rules or DNS rules using dynamic attributes filters:

Procedure

Step 1 Log in to Security Cloud Control as a user with the Super Admin role.

Click **Policies** > **Firewall Threat Defense** > **Access Control heading** > **Access Control**.

Step 2 Click **Firewall**.

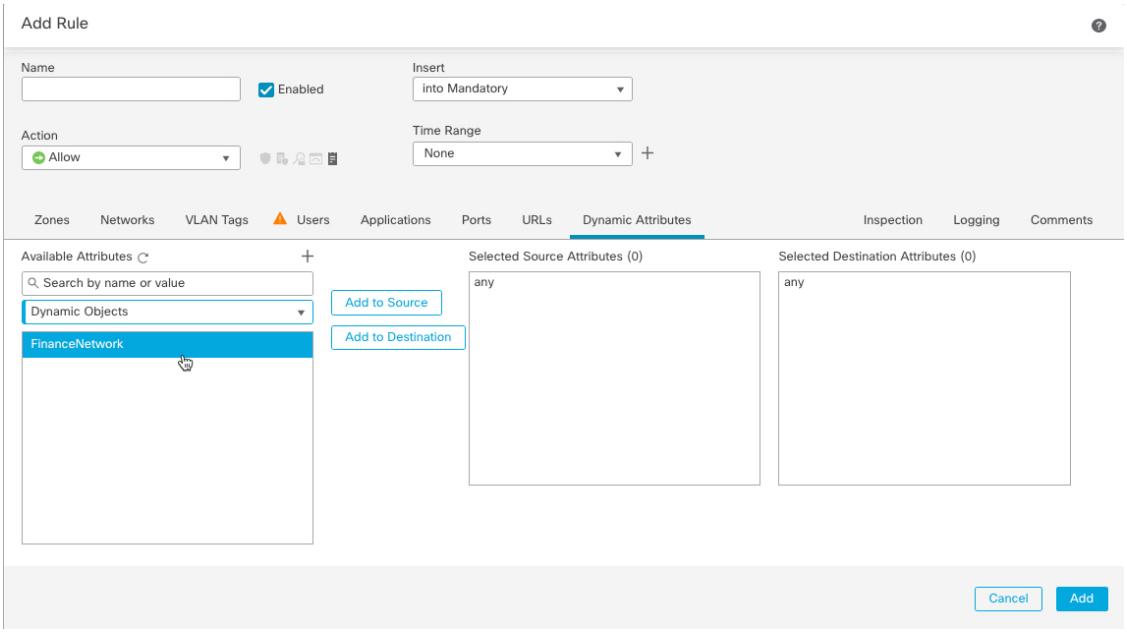
Step 3 Click **Edit** ( ) next to an access control policy.

Step 4 Click **Add Rule**.

Step 5 Click the **Dynamic Attributes** tab.

Step 6 In the Available Attributes section, from the list, click **Dynamic Objects**.

The following figure shows an example.



This example shows a dynamic object named `APIC Dynamic Attribute` that corresponds to the dynamic attribute filter created in the dynamic attributes connector.

Step 7 Add the desired object to source or destination attributes.

Step 8 Add other conditions to the rule if desired.

See Dynamic attributes rule conditions, on page 2681.

Use dynamic objects in DNS policies

Use dynamic objects in DNS rules to enable dynamic filtering capabilities based on attributes from the dynamic attributes connector.

The dynamic attributes connector enables you to configure dynamic filters seen in the Cloud-Delivered Firewall Management Center as dynamic objects, in DNS rules. For information about DNS policies, see DNS Policies for Security Intelligence, on page 2557.

A dynamic object is automatically pushed from the dynamic attributes connector to the Cloud-Delivered Firewall Management Center after you create connectors and save a dynamic attributes filter on the connector.

You can use these dynamic objects on the DNS rule's Dynamic Attributes tab page, similarly to the way you use Security Group Tags (SGTs). You can add dynamic objects as source or destination attributes, except for endpoint device type objects, which are source only.

Procedure

Step 1 Click **Policies** > **Access Control heading** > **DNS** and create or edit a DNS policy.

Step 2 Add or edit a rule.

Step 3 Click the **Dynamic Attributes** tab.

Step 4 In the **Dynamic Attributes** list, select the objects you want to use, then add them to the source or destination lists as appropriate.

Initially, all security group and dynamic objects are listed, by you can uncheck the Security Group option to see dynamic objects only.

Step 5 On the **DNS** tab, select the appropriate list or feed to match the DNS requests you are targeting.

Step 6 Add other conditions to the rule if desired and set the action.

Step 7 Click **Save**.

Troubleshoot the Dynamic Attributes Connector

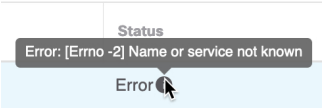
How to troubleshoot issues with the dynamic attributes connector, including using provided tools.

Troubleshoot error messages

This reference provides error messages and troubleshooting solutions for dynamic attributes connectors and adapters in the dynamic attributes connector.

Problem: name or service not known error

This error is displayed as a tooltip when you hover the mouse over an error condition on a connector. An example follows; yours might look different.

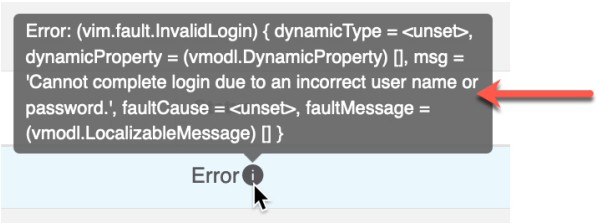


Solution: Edit the connector and check for:

- A trailing slash on a host name
- Verify the password is correct

Problem: incorrect username or password

This error is displayed as a tooltip when you hover the mouse over an error condition on a connector.



Solution: Edit the connector and change the user name or password.

Get your tenant ID

This task provides your tenant ID to Cisco TAC for troubleshooting assistance with the dynamic attributes connector.

If you require assistance with the dynamic attributes connector, you must provide your tenant ID to Cisco TAC so we can look at your logs.

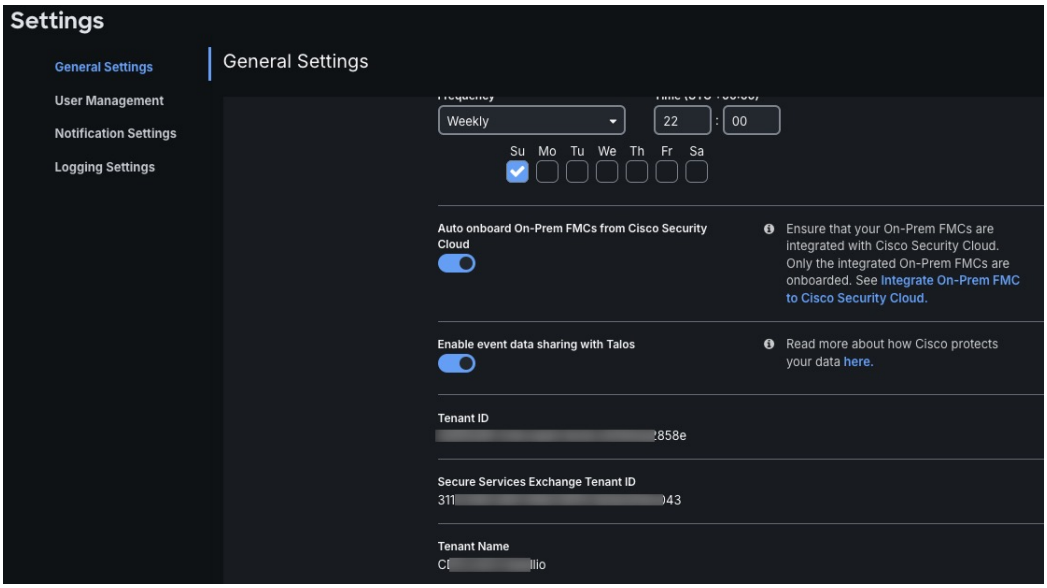
Procedure

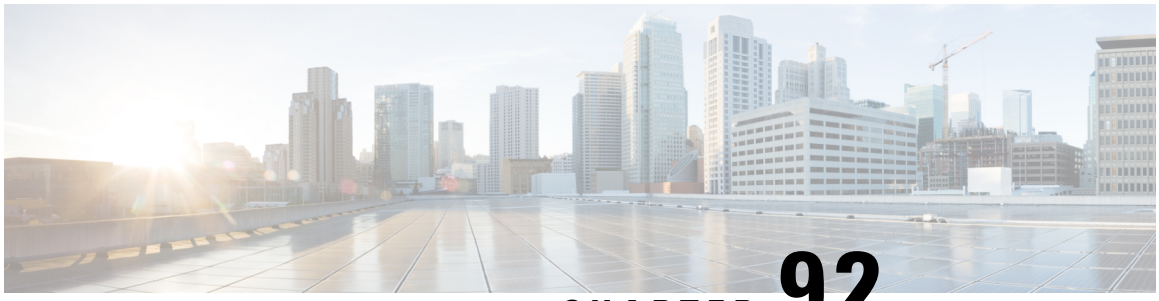
Step 1 Log in to Security Cloud Control.

Step 2 Click **Administration** > **General Settings**.

Step 3 Copy your tenant ID to the clipboard to provide to Cisco TAC.

A sample follows.


Settings
General Settings
User Management
Notification Settings
Logging Settings
General Settings
Weekly
22
00
Su Mo Tu We Th Fr Sa
Auto onboard On-Prem FMCs from Cisco Security Cloud
Ensure that your On-Prem FMCs are integrated with Cisco Security Cloud. Only the integrated On-Prem FMCs are onboarded. See Integrate On-Prem FMC to Cisco Security Cloud.
Enable event data sharing with Talos
Read more about how Cisco protects your data here.
Tenant ID
:858e
Secure Services Exchange Tenant ID
311)43
Tenant Name
C llio



CHAPTER 92

Advanced Settings for Access Control

Advanced access control policy settings typically require little or no modification. The default settings are appropriate for most deployments.

- Requirements and prerequisites for advanced settings, on page 2687
- Configure advanced settings for the access control policy, on page 2688

Requirements and prerequisites for advanced settings

This reference provides the requirements and prerequisites needed to configure advanced settings for access control policy, including model support, domain compatibility, and user role permissions.

Model support

Any

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin
- You can define custom user roles to differentiate between the intrusion configuration in access control policy and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. The existing pre-defined user roles that included the Modify Access Control Policy permission support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions. The granular permissions are:
  - **Policies** > **Access Control heading** > **Access Control** and choose the **Access Control Policy** > **Modify Access Control Policy** > **Modify Threat Configuration** allows the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for network analysis and intrusion policies, the configuration of the Security Intelligence policy for the access

control policy, and intrusion actions in the policy default action. If a user has this option only, the user can modify no other part of the policy or rule.

- **Modify Remaining Access Control Policy Configuration** controls the ability to edit all other aspects of the policy.

Configure advanced settings for the access control policy

Configure advanced settings to customize access control policy behavior for specific security requirements and performance optimization.

Advanced access control policy settings typically require little or no modification. The default settings are appropriate for most deployments. Note that many of the advanced preprocessing and performance options in access control policies may be modified by rule updates as described in *Update Intrusion Rules* in the *Cisco Secure Firewall Management Center Administration Guide* .

Caution See Configurations that restart the snort process when deployed or activated, on page 198 for a list of advanced setting modifications that restart the Snort process, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. Refer to Snort restart traffic behavior, on page 196.

Before you begin

Inheriting settings from a parent policy:

- If the access control policy has a base policy, you can elect to inherit settings from the base policy. Select **Inherit from (*base policy*)** for each setting group where you want to use the parent policy's settings. If you are allowed to configure unique settings for the policy, you must deselect the option to make your edits.
- If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. In this case, inheritance has been configured so that these settings are locked. The settings are read-only.

Follow these steps to configure advanced settings for the access control policy:

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Create or edit an access control policy.

Step 3 Select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 For each settings group, click **Edit** ( ) and configure the settings as needed.

For each feature group, a separate dialog box is opened where you can make your changes. Click **OK** to save any changes.

- **General Settings**—These settings apply broadly to the policy, including URL filtering options. For more information, see General settings, on page 2691.
- **Identity Policy Settings**—Select the policy to be used to implement user identity discovery. You must implement an identity policy to get user or user group information in connection events, or to write access control rules based on users or groups. For more information, see Identity policies, on page 2505.
- **Decryption Policy Settings**—Select the policy to be used when decrypting connections. You must decrypt traffic to apply inspection to encrypted connections.
- **TLS Server Identity Discovery**—Whether to allow the firewall to extract certificate details such as Common Name (CN), Organization, or Subject Alternative Names (SANs) even when TLS 1.3 encryption would normally hide them. This improves policy accuracy without requiring a decryption rule; the original client connection remains encrypted. For more information, see TLS server identity discovery, on page 2692.
- **Prefilter Policy Settings**—Select the policy to be used for statically offloading large flows, implementing early connection blocks, or rezoning plain-text tunnel traffic.
- **Network Analysis and Intrusion Policies**—Advanced network analysis and intrusion policy settings allow you to:
  - Specify the intrusion policy and associated variable set that are used to inspect packets that must pass before the system can determine exactly how to inspect that traffic.
  - Change the access control policy's default network analysis policy, which governs many preprocessing options.
  - Use custom network analysis rules and network analysis policies to tailor preprocessing options to specific security zones, networks, and VLANs.
- **Threat Defense Service Policy**—Use the service policy to apply services to specific traffic classes. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. The service policy rules are applied after the access control rules. For more information, see Service Policies, on page 2707.
- **Files and Malware Settings**—Tune file and malware inspection performance and storage, on page 2615 provides information on performance options for file control and malware defense.
- **Threat Detection**—Configure the portscan detector to detect and prevent portscan activity in all types of traffic to protect networks from eventual attacks. Portscan traffic can be detected efficiently in both allowed and denied traffic..
- **Elephant Flow Settings**—Elephant flows are large, long duration, and fast flows that can cause duress for Snort cores. There are two actions that you can apply on elephant flows to reduce system stress, CPU hogging, packet drops, and so on. These actions are:
  - Bypass any or all applications—This action bypasses flow from Snort inspection.
  - Throttle—This action applies dynamic rate limit policy (10% reduction) on elephant flows.

  For more information, see Configure elephant flow detection, on page 2736.
- **Intelligent Application Bypass Settings**—(Use **Elephant Flow Settings** instead of this option.) Intelligent application bypass (IAB) is an expert-level configuration that specifies applications to bypass or test for bypass if traffic exceeds a combination of inspection performance and flow thresholds.

IAB settings are applicable for Snort2 devices or pre 7.2.0 Snort3 devices. For more information, see Intelligent application bypass, on page 2693.

- **Transport/Network Layer Preprocessor Settings**—Advanced transport and network preprocessor settings apply globally to all networks, zones, and VLANs where you deploy your access control policy. For more information, see Transport/network layer preprocessor settings, on page 2699.
- **Detection Enhancement Settings**—Detection enhancement settings determine whether adaptive profiles are used for application detection and intrusion rules in the access control policy. Typically, the system uses the static settings in your network analysis policy to preprocess and analyze traffic. With adaptive profile updates, the system can adapt processing behavior using host information either detected by network discovery or imported from a third party. For more information, see Detection enhancement settings, on page 2700.
- **Performance Settings**—Settings for improving the performance of your system as it analyzes traffic for attempted intrusions. These settings are very advanced and should be left as defaults by most users. See Performance settings, on page 2702.
- **Latency-Based Performance Settings**—Settings specific to latency-based performance. These settings are very advanced and should be left as defaults by most users. See Latency-based performance settings, on page 2704.
- **Shadow Traffic**—The Shadow traffic dashboard enhances the visibility of traffic originating from unsanctioned privacy technologies. This type of traffic is specifically designed to evade traditional network monitoring and analysis by advanced firewalls. There is also a shadow traffic type attribute added to connection and unified events. You can disable this option if you do not need visibility on traffic that might contain unsanctioned content. For more information, see *The Shadow Traffic Widget* in the *Cisco Secure Firewall Management Center Administration Guide*
- **Advanced Logging**—Enable this feature to enrich connection logs with application data and forward the generated logs to the syslog destinations. Application logging leverages existing deep packet inspection capabilities to extract application data and enables you to enhance network monitoring and gain deeper insights into network traffic. This feature applies to Snort 3 Firewall Threat Defense devices.

  For more information about the application logging, see *.Application-Aware Event Logging* in the *Cisco Secure Firewall Management Center Administration Guide*

  Note
  Application logging can cause performance drop within network if used without filters configured in the access control rule. Filter specific traffic types using the access control rules to reduce the volume of logged traffic.

Step 5 Click **Save** to save the changes to the access control policy.

- To get back to the rules list, click **Access Control** in the packet flow line.
- You must deploy the configuration to apply changes to the managed devices.

The advanced settings for the access control policy are configured and saved. The changes will take effect after deployment to the managed devices.

General settings

General settings are advanced configuration options that

- control URL logging and character limits for connection events,
- manage user bypass timeouts for interactive blocking,
- determine URL cache lookup behavior and traffic inspection settings, and
- enable threat intelligence and DNS reputation enforcement features.

General advanced settings configuration options

These are the general advanced settings that you can configure for an access control policy:

- **Maximum URL characters to store in connection events**—Sets the maximum character length of each URL requested by users in end-of-connection events. Disabling or limiting the number of stored URL characters might improve system performance. The default is 1024. The range is 0 to 4096.

  Set the length to 0 to disable URL logging. Storing zero characters does not affect URL filtering. The system filters traffic based on requested URLs even though the system does not record them.

- **Allow an Interactive Block to bypass blocking for (seconds)**—Sets the time allowed for browsing after the user bypasses a URL filtering block. After the timeout expires, the user must bypass the block again. The default is 600 seconds (10 minutes). The range is 0 to 31536000 (8760 hours).

  Setting this value to 0 means the interactive block response is displayed once and the user bypass never expires.

- **Retry URL cache miss lookup**—This setting determines what the system does when it needs to look up a URL's category and reputation in the cloud.

  The first time the system encounters a URL that does not have a locally stored category and reputation, it looks up that URL in the cloud and adds the result to the local data store. This allows faster processing of that URL in the future.

  By default, this setting is enabled. The system momentarily delays the traffic while it checks the cloud for the URL's reputation and category, and uses the cloud verdict to handle the traffic.

  If you disable this setting, when the system encounters a URL that is not in its local cache, the traffic is immediately passed and handled according to the rules configured for uncategorized and reputationless traffic.

  In passive deployments, the system does not retry the lookup, as it cannot hold packets.

- **Enable reputation enforcement on DNS traffic**—Whether to have the system evaluate domain category and reputation early in URL transactions, when the browser looks up the domain name to get the IP address. Enable this option to improve URL filtering performance and efficacy. The default is enabled. For details and additional instructions, see DNS filtering, on page 2099 and subtopics.

- **Inspect traffic during policy apply**—Whether to inspect traffic when you deploy configuration changes unless specific configurations require restarting the Snort process. The default is enabled.

  When this option is enabled, resource demands could result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. The assigned device determines whether traffic drops during this interruption or passes without further inspection. See Snort restart scenarios, on page 194 for more information.

TLS server identity discovery

TLS server identity discovery is an access control policy advanced setting that enables early application detection and URL categorization for encrypted traffic.

TLS server identity discovery configuration

The latest version of the Transport Layer Security (TLS) protocol 1.3, defined by RFC 8446, is the preferred protocol for many web servers to provide secure communications. Because the TLS 1.3 protocol encrypts the server's certificate for additional security, and the certificate is needed to match application and URL filtering criteria in access control rules, the Firepower System provides a way to extract the server certificate *without* decrypting the entire packet.

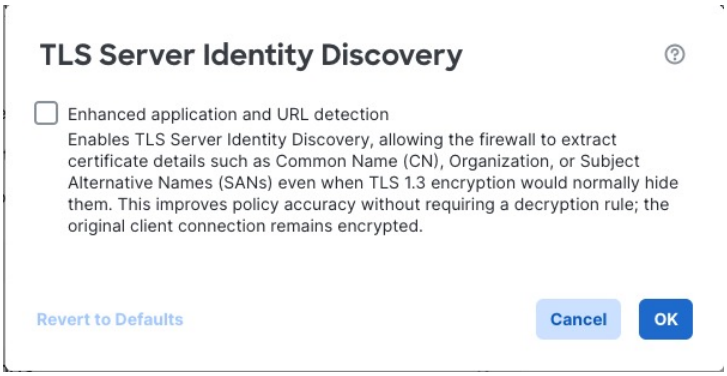
You can enable this feature, referred to as *TLS server identity discovery*, when you configure advanced settings for an access control policy. Certain features are not supported, such as STARTTLS traffic, the HTTP CONNECT method, and in a network where another device is already decrypting traffic.

If you enable this option, we recommend you also enable the decryption policy's advanced TLS adaptive server identity probe option as well. Together, these options enable more efficient decryption of TLS 1.3 traffic. For more information, see TLS 1.3 decryption best practices, on page 2184.

When a new connection starts that will be affected by TLS server identity discovery, the Firewall Threat Defense holds the original ClientHello packet to determine the identity of the server to which it connects before continuing. The Firewall Threat Defense device sends a specialized connection from the Firewall Threat Defense to the server. The server's response includes the server certificate, the specialized connection is terminated, and the original connection is evaluated as required by the access control policy.

TLS server identity discovery prioritizes the certificate's Common Name (CN) over the Server Name Indication (SNI).

To enable TLS server identity discovery, click the **Advanced** tab, click **Edit** (✎) for the setting, and select **Early application detection and URL categorization**.



We strongly recommend enabling it for any traffic you want to match on application or URL criteria, especially if you want to perform deep inspection of that traffic. An decryption policy is not required because *traffic is not decrypted* in the process of extracting the server certificate.

Note

- TLS server identity discovery *cannot* be used with any of the following:
  - STARTTLS traffic
  - The HTTP CONNECT method
  - Traffic that is already being decrypted by another device on the network
- Because the certificate is decrypted, TLS server identity discovery can reduce performance depending on the hardware platform.
- TLS server identity discovery is not supported in inline tap mode or passive mode deployments.
- Enabling TLS server identity discovery is supported on any Secure Firewall Threat Defense Virtual version deployed in AWS. TLS server identity is incompatible with certain deployment scenarios that involve Gateway Load Balancer (GWLB). Whether the device initiates a server certificate probe depends on its deployment method with GWLB and the traffic direction.
  - Geneve single-arm flow: The probe *is not initiated* for either east–west or north–south traffic. A connection event **PROBE_FLOW_DROP_BYPASS_PROXY** increments every time the device attempts to extract the server certificate.
  - Geneve dual-arm flow: The probe *is initiated only* for north–south traffic.
- TLS Server Identity Discovery also operates on TLS 1.2 sessions.

Intelligent application bypass

Intelligent application bypass (IAB) is a security feature that

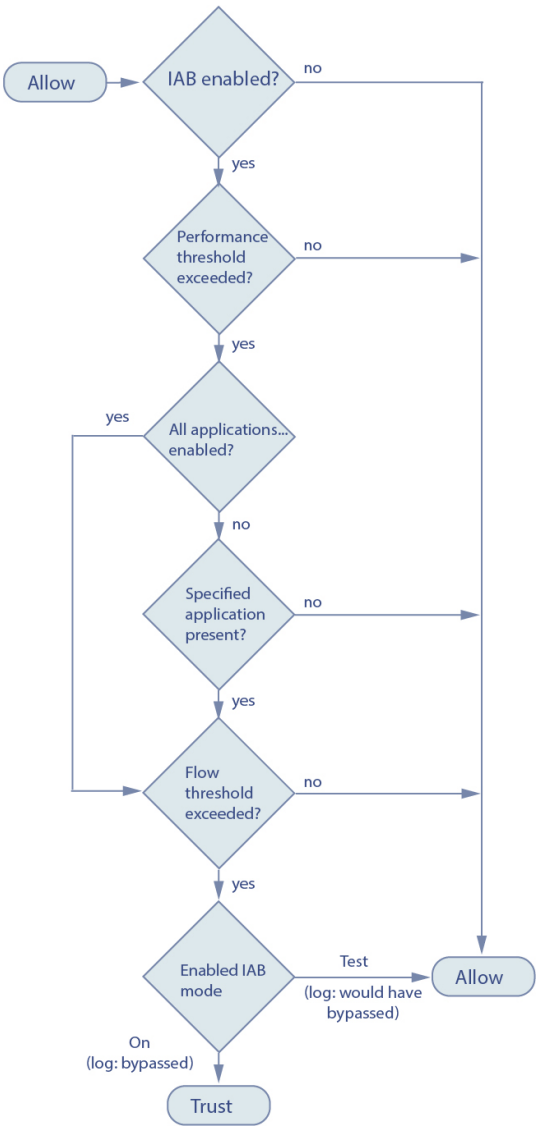
- identifies applications that you trust to traverse your network without further inspection if performance and flow thresholds are exceeded
- implements on traffic allowed by access control rules or the access control policy's default action, before the traffic is subject to deep inspection, and
- can be configured so that when an inspection performance threshold is exceeded, IAB trusts all traffic that exceeds any flow bypass threshold, regardless of the application type.

Additional information

For example, if a nightly backup significantly impacts system performance, you can configure thresholds that, if exceeded, trust traffic generated by your backup application.

A test mode allows you to determine whether thresholds are exceeded and, if so, to identify the application flows that would have been bypassed if you had actually enabled IAB (called *bypass mode*).

This graphic illustrates the IAB decision-making process:



Configure intelligent application bypass

This task enables intelligent application bypass (IAB) to improve system performance by bypassing inspection of specified traffic when performance thresholds are exceeded.

Not all deployments require IAB, and those that do might use it in a limited fashion. Do not enable IAB unless you have expert knowledge of your network traffic, especially application traffic, and system performance, including the causes of predictable performance issues. Before you run IAB in bypass mode, make sure that trusting the specified traffic does not expose you to risk.

Before you begin

IAB settings are applicable for Snort2 devices or pre 7.2.0 Snort3 devices. For Snort 3 devices, use elephant flow detection instead.

Follow these steps to configure intelligent application bypass:

Procedure

Step 1 In the access control policy editor, click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line. Then, click **Edit** (✎) next to **Intelligent Application Bypass Settings**.

Step 2 Set the enabled **State** for IAB.

Turn IAB **Off** or **On**, or enable IAB in **Test** mode.

In test mode, connection events and dashboards tell you what the system would have done if IAB had been on, but traffic is not impacted. Use test mode to check your configuration.

Step 3 Set the **Performance Sample Interval**.

The **Performance Sample Interval** specifies the time in seconds between IAB performance sampling scans, during which the system collects system performance metrics for comparison to IAB performance thresholds. The default is 5 seconds. The range is 1 to 1000 seconds.

Step 4 Select the **Bypassable Applications and Filters**.

Choose from:

- ***X* Applications/Filters**—Click the link and select the applications or application filters whose traffic you want to bypass. You can select by general attributes, specific applications, or both. For example, you could limit bypassable traffic to very low risk applications only.
- **All applications including unidentified applications**—Do not restrict bypass. When an inspection performance threshold is exceeded, IAB trusts all traffic that exceeds any flow bypass threshold, regardless of the application type. This is the default.

Step 5 Configure the performance and flow thresholds.

You must configure at least one **Inspection Performance Thresholds** and one **Flow Bypass Thresholds**. However, all have defaults and you do not need to change the settings if the defaults are appropriate for your network.

When a performance threshold is exceeded, the system examines flow thresholds and, if one threshold is exceeded, trusts the specified traffic. If you enable more than one of either, only one of each must be exceeded.

a) Click **Configure** under **Inspection Performance Thresholds** and configure the options.

Inspection performance thresholds provide intrusion inspection performance limits that, if exceeded, trigger the inspection of flow thresholds. Inspection performance thresholds set to 0 are ignored.

You can configure one or more of these thresholds:

- **Drop Percentage**—The average packets dropped as a percentage of total packets, when packets are dropped because of performance overloads caused by expensive intrusion rules, file policies, decompression, and so on. This does not refer to packets dropped by normal configurations such as intrusion rules.

  Note that specifying an integer greater than 1 activates IAB when the specified percentage of packets is dropped. When you specify 1, any percentage from 0 through 1 activates IAB. This allows a small number of packets to activate IAB.

  The default is 5%. The range is 0 to 100.

- **Processor Utilization Percentage**—Average percentage of processor resources used. The default is 95. The range is 0 to 100.
- **Packet Latency**—The average packet latency in microseconds. The default is 1000. The range is 0 to 1000000.
- **Flow Rate**—The rate at which the system processes flows, measured as the number of flows per second. Note that this option configures IAB to measure flow *rate*, not flow *count*. The default is 0. The range is 0 to 1000000.

b) Click **Configure** under **Flow Bypass Thresholds** and configure the options.

Flow bypass thresholds provide flow limits that, if exceeded, trigger IAB to trust bypassable application traffic in bypass mode or allow application traffic subject to further inspection in test mode. Flow bypass thresholds set to 0 are ignored.

You can configure one or more of these flow bypass thresholds:

- **Bytes per Flow**—The maximum number of kilobytes a flow can include. The default is 500000. The range is 0 to 2147483647.
- **Packets per Flow**—The maximum number of packets a flow can include. The default is 0. The range is 0 to 2147483647.
- **Flow Duration**—The maximum number of seconds a flow can remain open. The default is 0. The range is 0 to 2147483647.
- **Flow Velocity**—The maximum transfer rate in kilobytes per second. The default is 250000. The range is 0 to 2147483647.

Step 6 Click **OK** to save IAB settings.

Step 7 Click **Save** to save the policy.

Intelligent application bypass is configured for the access control policy with the specified performance and flow thresholds.

What to do next

- Deploy configuration changes.

IAB logging and analysis

IAB logging and analysis is a monitoring system that

- forces an end-of-connection event. This event logs both bypassed flows and flows that would have been bypassed, regardless of whether you have enabled connection logging
- provides connection events. These events indicate when flows are bypassed in bypass mode or would have been bypassed in test mode, and
- enables custom dashboard widgets and reports based on connection events. These widgets and reports display long-term statistics for bypassed and would-have-bypassed flows.

IAB connection events

IAB connection events provide specific information about bypassed and would-have-bypassed flows through these key fields:

- Action: When **Reason** includes `Intelligent App Bypass`:
  - **Allow**—indicates that the applied IAB configuration was in test mode and traffic for the application specified by **Application Protocol** remains available for inspection.
  - **Trust**—indicates that the applied IAB configuration was in bypass mode and traffic for the application specified by **Application Protocol** has been trusted to traverse the network without further inspection.
- Reason: `Intelligent App Bypass` indicates that IAB triggered the event in bypass or test mode.
- Application Protocol: This field displays the application protocol that triggered the event.

You can create custom analysis dashboard widgets to display long-term IAB statistics based on connection events. Specify these settings when creating the widget:

- **Preset**: None
- **Table**: Application Statistics
- **Field**: any
- **Aggregate**: either of:
  - IAB Bypassed Connections
  - IAB Would Bypass Connections
- **Filter**: any

You can create custom reports to display long-term IAB statistics based on connection events. Specify these settings when creating the report:

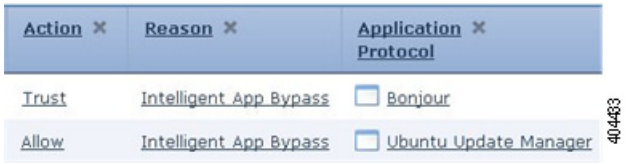
- **Table**: Application Statistics
- **Preset**: None
- **Filter**: any
- **X-Axis**: any
- **Y-Axis**: either of:
  - IAB Bypassed Connections
  - IAB Would Bypass Connections

Connection event examples

In the truncated graphic, some fields are omitted. The graphic shows the **Action**, **Reason**, and **Application Protocol** fields for two connection events resulting from different IAB settings in two separate access control policies.

For the first event, the `Trust` action indicates that IAB was enabled in bypass mode and Bonjour protocol traffic was trusted to pass without further inspection.

For the second event, the `Allow` action indicates that IAB was enabled in test mode, so Ubuntu Update Manager traffic was subject to further inspection but would have been bypassed if IAB had been in bypass mode.

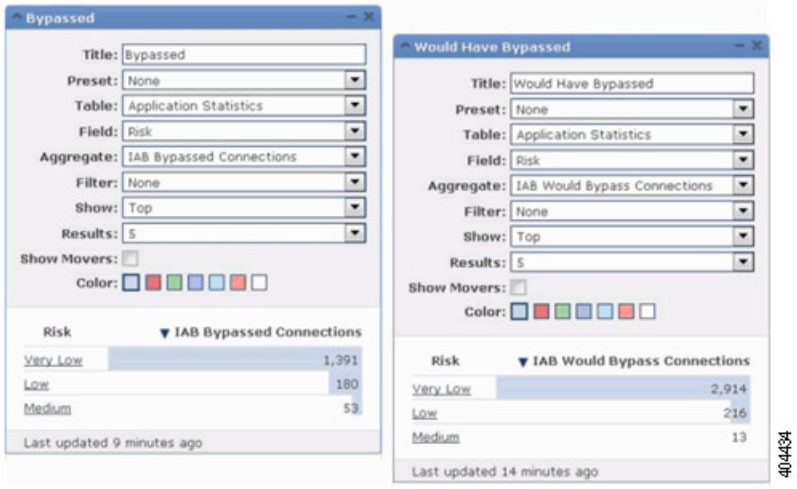


In the truncated graphic, some fields are omitted. The flow in the second event was both bypassed (**Action**: `Trust`; **Reason**: `Intelligent App Bypass`) and inspected by an intrusion rule (**Reason**: `Intrusion Monitor`). The `Intrusion Monitor` reason indicates that an intrusion rule set to **Generate Events** detected but did not block an exploit during the connection. In the example, this happened before the application was detected. After the application was detected, IAB recognized the application as bypassable and trusted the flow.



In the custom analysis dashboard widget examples:

- The *Bypassed* example shows statistics for application traffic bypassed because the applications were specified as bypassable and IAB was enabled in bypass mode in the deployed access control policy.
- The *Would Have Bypassed* example shows statistics for application traffic that would have been bypassed because the applications were specified as bypassable and IAB was enabled in test mode in the deployed access control policy.



The graphic shows two abbreviated report examples:

- The *Bypassed* example shows statistics for application traffic bypassed because the applications were specified as bypassable and IAB was enabled in bypass mode in the deployed access control policy.
- The *Would Have Bypassed* example shows statistics for application traffic that would have been bypassed because the applications were specified as bypassable and IAB was enabled in test mode in the deployed access control policy.



Transport/network layer preprocessor settings

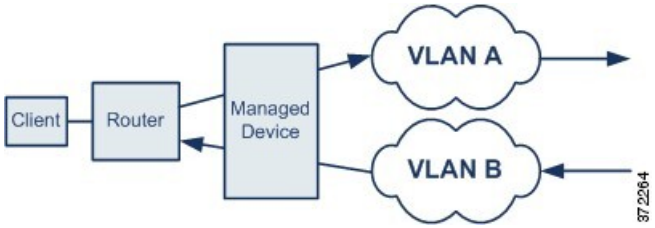
Transport/network layer preprocessor settings are access control policy configurations that

- apply globally to all networks, zones, and VLANs where you deploy your access control policy
- control how the system handles VLAN headers, active responses, and connection tracking, and
- affect traffic reassembly and rule processing for transport and network layer communications.

Configuration options

These settings control various aspects of transport and network layer processing:

- **Ignore the VLAN header when tracking connections**—Whether to ignore or include VLAN headers when identifying traffic. Different VLAN tags in traffic traveling in different directions for the same connection can affect traffic reassembly and rule processing. For example, traffic for the same connection could be transmitted over VLAN A and be received over VLAN B. Select this option if the device might see different VLANs for the same connection. This option is off by default.



- **Maximum Active Responses**—For a TCP connection that triggers a preprocessor/intrusion drop rule that is configured to provide an active response, the maximum number of active responses per TCP connection. When additional traffic occurs on a connection where an active response has been initiated, and the traffic occurs more than **Minimum Response Seconds** after a previous active response, the system sends another active response unless the specified maximum has been reached. A setting of 0 disables additional active responses triggered active response rules. The default is no limit. The range is 0 to 25.

Note You have to specifically configure drop rules to provide active responses. For TCP connections, the active response is a RESET packet. For UDP connections, the system sends an ICMP unreachable packet to the source of the connection.

- **Minimum Response Seconds**—Until **Maximum Active Responses** occur, specifies the number of seconds to wait before any additional traffic on a connection where the system has initiated an active response results in a subsequent active response. The default is no limit. The range is 1 to 300.
- **Session Termination Logging Threshold**—Do not modify this option unless instructed to do so by Cisco Technical Support.

  Support might ask you during a troubleshooting call to configure your system to log a message when an individual connection exceeds the specified threshold. Changing the setting for this option will affect performance and should be done only with Support guidance. This option specifies the number of bytes that result in a logged message when the session terminates and the specified number was exceeded. The upper limit is 1GB.

Detection enhancement settings

Detection enhancement settings are adaptive profile configurations that

- determine whether adaptive profiles are used for application detection and intrusion rules in the access control policy
- enable the system to adapt processing behavior using host information either detected by network discovery or imported from a third party, and
- provide control over profile update frequency and network scope for performance optimization.

Detection enhancement setting options

Typically, the system uses the static settings in your network analysis policy to preprocess and analyze traffic. With adaptive profile updates enabled, the system adapts processing behavior using host information detected by network discovery or imported from a third party.

Note To enable adaptive profiles in Snort 3, you must select both the **Enable** and **Enable Profile Updates** options.

The detection enhancement settings include these options:

- **Enable**—You must enable adaptive profiling (its default state) for access control rules to perform application and file control, including malware protection (AMP), and for intrusion rules to use service metadata.
- **Enable Profile Updates**—Profile updates, like the target-based profiles you can configure manually in a network analysis policy, help to defragment IP packets and reassemble streams in the same way as the operating system on the target host. The intrusion rules engine then analyzes the data in the same format as that used by the destination host. Profile updates also compare metadata in an intrusion rule to host information to determine whether a rule should apply for a particular host. For more information, see:

- Adaptive profile updates, on page 2701
- Adaptive profile updates and recommended rules, on page 2702

- **Adaptive Profiles – Attribute Update Interval**—When profile updates are enabled, you can control how frequently in minutes network map data is synced from the management center to its managed devices. The system uses the data to determine what profiles should be used when processing traffic. Increasing the value for this option can improve performance in a large network.
- **Adaptive Profiles – Networks**—Optionally, when profile updates are enabled, you can improve performance by constraining profile updates to a comma-separated list of IP addresses, address blocks, and network variables. If you use a network variable, the system uses the variable's value in the variable set linked to the default intrusion policy for your access control policy. For example, you could enter 192.168.1.101, 192.168.4.0/24, or $HOME_NET. IPv4 and IPv6 are supported.

  The default value (0.0.0.0/0) applies adaptive profile updates to all networks.

Adaptive profile updates

An adaptive profile update is a network analysis feature that

- adapts processing behavior using host information either detected by network discovery or imported from a third party
- switches to the appropriate operating system profile based on the operating system in the host profile for the target host, and
- helps to defragment IP packets and reassemble streams in the same way as the operating system on the target host.

Adaptive profile operation

Typically, the system uses the static settings in your network analysis policy to preprocess and analyze traffic. With adaptive profile updates, the system can adapt processing behavior using host information either detected by network discovery or imported from a third party.

Profile updates, like the target-based profiles you can configure manually in a network analysis policy, help to defragment IP packets and reassemble streams in the same way as the operating system on the target host. The intrusion rules engine then analyzes the data in the same format as that used by the destination host.

Manually configured target-based profiles apply either the default operating system profile you select, or profiles you bind to specific hosts. Profile updates, however, switch to the appropriate operating system profile based on the operating system in the host profile for the target host.

Adaptive profile update scenario

Consider a scenario where you configure profile updates for the 10.6.0.0/16 subnet and set the default IP Defragmentation target-based policy to Linux. The Cloud-Delivered Firewall Management Center where you configure the settings has a network map that includes the 10.6.0.0/16 subnet.

- When the system detects traffic from Host A, which is not in the 10.6.0.0/16 subnet, it uses the Linux target-based policy to reassemble IP fragments.

- When the system detects traffic from Host B, which is in the 10.6.0.0/16 subnet, it retrieves Host B's operating system data from the network map. The system uses a profile based on that operating system to defragment the traffic destined for Host B.

Adaptive profile updates and recommended rules

Adaptive profile updates and recommended rules are security features that

- compare metadata in a rule to host information to determine whether a rule should apply for a particular host
- provide different approaches to rule optimization, with adaptive profiles applying rules dynamically on a packet-by-packet basis while recommended rules provide suggestions for rule state changes, and
- can be used in combination to ensure appropriate rule enablement and efficient rule application for specific traffic.

Feature characteristics and differences

The adaptive profile updates feature is an advanced setting in an access control policy that applies globally to all intrusion policies invoked by that access control policy. The Cisco recommended rules feature applies to the individual intrusion policy where you configure it.

Like recommended rules, profile updates compare metadata in a rule to host information to determine whether a rule should apply for a particular host. However, while recommended rules provide recommendations for enabling or disabling rules using that information, profile updates use the information to apply specific rules to specific traffic.

This table compares the key differences between adaptive profile updates and recommended rules:

Recommended rules require your interaction to implement suggested changes to rule states. Profile updates, on the other hand, do not modify intrusion policies. Treatment of rules based on profile updates happens on a packet-by-packet basis.

Additionally, recommended rules can result in enabling disabled rules. Profile updates, in contrast, only affect the application of rules that are already enabled in intrusion policies. Profile updates never change the rule state.

You can use profile updates and recommended rules in combination. Profile updates use the rule state for a rule when your intrusion policy is deployed to determine whether to include it as a candidate for applying, and your choices to accept or decline recommendations are reflected in that rule state. You can use both features to ensure that you have enabled or disabled the most appropriate rules for each network you monitor, and then to apply enabled rules most efficiently for specific traffic.

Performance settings

Performance settings are access control policy advanced settings that

- tune the performance of your system as it analyzes traffic for attempted intrusions
- apply to Snort and are relevant only if you apply intrusion policies in rules or as the default action, and
- should be changed only if you are a Snort intrusion rule expert, or at the direction of Cisco Technical Support.

Performance setting categories

These settings are organized into separate tabs for each group of settings:

- **Pattern Matching Limits**: The maximum number of events to queue and TCP payload detection options
- **Performance Statistics**: Time range and packet count settings for performance sampling
- **Regular Expression Limits**: PCRE matching and recursion limits
- **Intrusion Event Logging Limits**: Event queue and logging configuration

Pattern Matching Limits include these options:

- **Maximum Pattern States to Analyze Per Packet**: The maximum number of events to queue. The default is 5.
- **Disable Content Checks on Traffic Subject to Future Reassembly**: Whether to detect TCP payload before reassembly. It includes inspection of packets before and after stream reassembly. This process requires more processing overhead and may decrease performance. If the option is not selected, the TCP payload is detected after reassembly. The default is off.

Performance Statistics settings control when the system updates performance statistics. When the number of seconds specified elapses between performance statistics updates, the system verifies it has analyzed the specified number of packets. If it has, the system updates performance statistics. Otherwise, the system waits until it analyzes the specified number of packets.

- **Sample time (seconds)**: The time range for taking a performance sample. The default is 300 seconds.

  Configuring a very low value (for example, 1 second) for the sample time can cause a huge impact on the device; the performance statistics logged on the device can cause disk space issues and affect the operation of the device. Hence, we recommend you do not configure a very low value.

- **Minimum number of packets**: How many packets to consider the minimum for a valid performance statistic. The default is 0.
- **Troubleshooting Options**:
  - **Log Session/Protocol Distribution**: Support might ask you during a troubleshooting call to enable this option to configure the system to calculate the performance statistics only when the Snort process is shut down or restarted.
  - **Summary**: Enable this option only if instructed by Cisco Technical Support.

Regular Expression Limits ensure a minimum level of performance. The default Perl compatible regular expression (PCRE) limits ensure a minimum level of performance. Overriding these limits could increase security but could also significantly impact performance by permitting packet evaluation against inefficient regular expressions.

- **Match Limit State**: The limit for matching regular expressions. You can select **Default Value**, which is 3500, **Unlimited**, or **Custom**. If you select custom, specify the number of times to attempt to match a pattern defined in a PCRE regular expression in **Match Limit**. Specify 0 to completely disable PCRE match evaluations.
- **Match Recursion Limit State**: The limit for matching regular expression recursions. You can select **Default Value**, which is 3500, **Unlimited**, or **Custom**. If you select custom, specify the number of

recursions when evaluating a PCRE regular expression against the packet payload in **Match Recursion Limit**. Specify 0 to completely disable PCRE recursions.

Note For a custom match recursion limit to be meaningful, it must be smaller than the match limit.

Intrusion Event Logging Limits control event queuing and reporting. When the intrusion rules engine evaluates traffic against rules, it places the events generated for a given packet or packet stream in an event queue, then reports the top events in the queue to the user interface. When configuring the intrusion event logging limits, you can specify how many events can be placed in the queue and how many are logged, and select the criteria for determining event order within the queue.

- **Maximum Events Stored Per Packet**: The maximum number of events that can be stored for a given packet or packet stream. The default is 8.
- **Maximum Events Logged Per Packet**: The number of events logged for a given packet or packet stream. This cannot exceed the **Maximum Events Stored Per Packet** value. The default is 5.
- **Prioritize Events Logging By**: The value used to determine event ordering within the event queue. The highest ordered event is reported through the user interface. You can select:
  - **Content Length** (the default), which orders events by the longest identified content match. When events are ordered by content length, rule events always take precedence over decoder and preprocessor events.
  - **Priority**, which orders events in the queue by the event priority.

Latency-based performance settings

Latency-based performance settings are policy configurations that

- use thresholding to manage packet and rule processing performance,
- apply to Snort and are relevant only if you apply intrusion policies in rules or as the default action, and
- are automatically populated by the latest deployed intrusion rule update by default.

Performance setting behavior

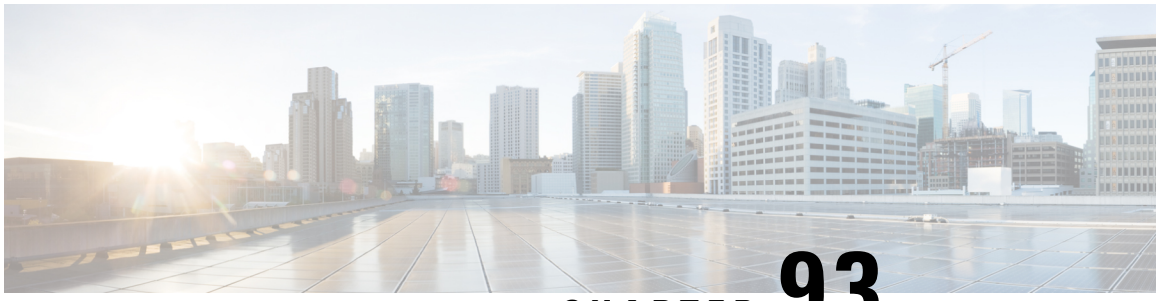
By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and we recommend that you do not change the default. Change these settings only if you are a Snort intrusion rule expert, or at the direction of Cisco Technical Support.

The latency settings that are applied depend on the security level of the network analysis policy (NAP) associated with the access control policy. Generally, this is the default NAP policy. However, if custom network analysis rules are configured, and if any of these specify a NAP policy that is more secure than the default NAP policy, then latency settings are based on the most secure NAP policy among the custom rules. If the default NAP policy or any custom rules invoke a custom NAP policy, then the security level used in the evaluation is the system-provided base policy on which each custom NAP policy is based.

The above is true regardless of whether the effective threshold and/or network analysis configurations are inherited or configured directly in the policy.

Use these settings to tune the latency-based performance of your system:

- **Apply Settings From**: Whether to apply latency-based performance settings from an **Installed Rule Update**, the default, or from your **Custom** settings.
- **Packet Handling**: Packet latency thresholding measures the total elapsed time taken to process a packet by applicable decoders, preprocessors, and rules, and ceases inspection of the packet if the processing time exceeds a configurable threshold. By default, the latency-based performance setting for packet handling is disabled. You may choose to enable it. However, Cisco recommends that you do not change the default value for the threshold setting. Select **Enabled** to turn it on. If you selected **Custom**, also enter the **Threshold** time in microseconds for when inspection of a packet should cease. The default is 256.
- **Rule Handling**: Rule latency thresholding measures the elapsed time each rule takes to process an individual packet, suspends the violating rule along with a group of related rules for a specified time if the processing time exceeds the rule latency threshold a configurable consecutive number of times, and restores the rules when the suspension expires. You can configure these options only if you select **Custom**.
  - **Enabled**: This option is selected automatically if you select custom. If you do not want use this feature, deselect the option. All other settings on this tab require the feature to be enabled.
  - **Threshold (microseconds)**: Specifies the time in microseconds that rules should not exceed when examining a packet. The default is 512.
  - **Consecutive Threshold Violations Before Suspending Rule**: Specifies the consecutive number of times rules can take longer than the time set for threshold to inspect packets before rules are suspended. The default is 3.
  - **Suspension Time**: How long a suspended rule should be suspended. The default is 10.



CHAPTER 93

Service Policies

You can use Threat Defense Service Policies to apply services to specific traffic classes. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. A service policy consists of multiple actions or rules applied to an interface or applied globally.

- Threat defense service policies, on page 2707
- Requirements and prerequisites for service policies, on page 2709
- Guidelines and limitations for service policies, on page 2709
- Configure threat defense service policies, on page 2710
- Monitor service policies, on page 2719
- Examples for service policy rules, on page 2720

Threat defense service policies

Use Threat Defense service policies to apply services to specific traffic classes. Service policies let you apply different services to different connections entering the device or interface.

Traffic class types

A traffic class is a combination of the interface and an extended access control list (ACL). The ACL "allow" rules determine which connections are part of the class. Any "denied" traffic in the ACL simply does not have the service applied to it: these connections are not actually dropped. You can use IP addresses and TCP/UCP ports to identify matching connections as precisely as you require.

There are two types of traffic class:

- **Interface-based rules**—If you specify a security zone or interface group in a service policy rule, the rule applies to the ACL "allowed" traffic that goes through any interface that is part of the interface objects.

  Interface-based rules that are applied to the ingress interface always take precedence over global rules for any feature. If an ingress interface-based rule applies to a connection, the system ignores any matching global rule. If no ingress interface or global rule applies, then an interface service rule on the egress interface is applied.

- **Global rules**—These rules apply to all interfaces. If an interface-based rule does not apply to a connection, the global rules are checked and applied to any connections that the ACL "allows." If none apply, then the connections proceed without any services applied.

A given connection can match only one traffic class, either interface-based or global, for a given feature. There should be at most one rule for a given interface object/traffic flow combination.

Service policy rules are applied after access control rules. These services are configured only for connections you are allowing.

Service policy relationships with FlexConfig and other features

Connection-related service policy features are managed separately from other service-rule features, so overlapping traffic classes generally aren't an issue, but they should still be configured carefully.

Service policy feature interactions

Because connection-related service policy features are treated as a separate feature group from other service-rule implemented features, you should not run into problems with overlapping traffic classes. However, please be mindful when configuring these features:

- QoS policy rules: QoS policy rules are implemented using the service policy CLI. These rules are applied before connection-based service policy rules. However, both QoS and connection settings can be applied to the same or overlapping traffic classes.
- FlexConfig policies: You can use FlexConfig policies to implement custom application inspections and NetFlow. Use the **show running-config** command to examine the CLI that already configures service rules, including the **policy-map**, **class-map**, and **service-policy** commands. Netflow and application inspection are compatible with QoS and connection settings, but you need to understand the existing configuration before implementing FlexConfig. Connection settings are applied before application inspections and Netflow.



Note Traffic classes that are created from the Threat Defense service policy are named **class_map_***ACLname*, where *ACLname* is the name of the extended ACL object used in the service policy rule.

Connection settings

Connection settings refer to a variety of features that manage connections through a device.

Connection setting types

Connection settings include these types:

- **Global timeouts for various protocols**—All global timeouts have default values. Change them only if you experience premature connection loss. You configure global timeouts in the Threat Defense platform policy. Select **Devices** > **Platform Settings**.
- **Connection timeouts per traffic class**—You can override the global timeouts for specific types of traffic using service policies. All traffic class timeouts have default values, so you do not have to set them.
- **Connection limits and TCP Intercept**—By default, there are no limits on how many connections can go through (or to) the Firewall Threat Defense. Use service policy rules to set limits on particular traffic classes. This protects servers from denial of service attacks Particularly, you can set limits on embryonic connections (those that have not finished the TCP handshake), which protects against SYN flooding

attacks. When embryonic limits are exceeded, the TCP Intercept component activates to proxy connections and throttle attacks.

- **Dead connection detection (DCD)**—Enable dead connection detection if you have persistent connections that are valid but often idle. These connections may get closed when they exceed idle timeout settings. Dead connection detection identifies idle but valid connections and keeps them alive by resetting their idle timers. Whenever idle times are exceeded, DCD probes both sides of the connection to see if both sides agree that the connection is valid.. The **show service-policy** command output includes counters to show the amount of activity from DCD. You can use the **show conn detail** command to get information about the initiator and responder and how often each has sent probes.
- **TCP sequence randomization**—Each TCP connection has two initial sequence numbers (ISN): one generated by the client and one generated by the server. By default, the Firewall Threat Defense randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions. Randomization prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session. However, TCP sequence randomization effectively breaks TCP SACK (Selective Acknowledgement), as the sequence numbers the client sees are different from what the server sees. You can disable randomization per traffic class if desired.
- **TCP normalization**—The TCP Normalizer protects against abnormal packets. You can configure how some types of packet abnormalities are handled by traffic class. You can configure TCP Normalization using the FlexConfig policy.
- **TCP state bypass**—You can bypass TCP state checking if you use asymmetrical routing in your network.

Requirements and prerequisites for service policies

This topic outlines the model support, domain compatibility, and user role requirements for configuring service policies.

Model Support

Firewall Threat Defense

Supported domains

Any

User roles

Admin

Access Admin

Network Admin

Guidelines and limitations for service policies

Service policies apply to routed or switch interfaces only, in either routed or transparent mode. They do not apply to inline set or passive interfaces.

You can have at most 25 traffic classes for a given interface or the global policy. Specifically, this means that you cannot have more than 25 service policy rules for the global policy for a given security zone or interface group. However, for interfaces, because the same interface can appear in both a security zone and interface group, be aware that the actual limitation is based on the interfaces, and not the zone/group. Thus, you might be prevented from having 25 rules per zone/group based on the membership of your zones/groups.

You can have at most one rule for a given interface object/traffic flow combination.

When you make service policy changes to the configuration, all new connections use the new service policy. Existing connections continue to use the policy that was configured at the time of the connection establishment. If you want all connections to immediately use the new policy, you need to disconnect the current connections so they can reconnect using the new policy. Use an SSH or Console CLI session to enter the **clear conn** or **clear local-host** commands.

QoS rules and service policy rules (in the access control policy) are translated to commands on the firewall. This includes creating policy maps. There can be only one policy map per interface. The system automatically combines QoS and service policy configurations in a single policy map per interface.

However, you can also create policy maps manually using FlexConfig. For lower-end firewalls, such as the Secure Firewall 3100, you can create up to 64 policy maps in the configuration. For more powerful firewalls, such as the Firepower 4100, you can create up to 128. Be careful not to exceed this limit, or deployment will fail. You can view the device configuration from the CLI or use the deployment transcript to see how rules are translated to commands.

Configure threat defense service policies

You can use Threat Defense service policies to apply services to specific traffic classes. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications.

A service policy consists of multiple actions or rules applied to an interface or applied globally. Use Threat Defense service policies to apply services to specific traffic classes.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**, and click **Edit** (✎) for the access control policy whose Threat Defense service policy you want to edit.

Step 2 Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 3 Click **Edit** (✎) in the **Threat Defense Service Policy** group.

A dialog box opens that shows the existing policy. The policy consists of an ordered list of rules, separated between global rules (which apply to all interfaces) and interface-based rules. The table shows the interface object and extended access control list name (which combined defines the traffic class for the rule), and the services applied.

Step 4 Do any of the following:

- Click **Add Rule** to create a new rule. See Configure a service policy rule, on page 2711.
- Click **Edit** (✎) to edit an existing rule. See Configure a service policy rule, on page 2711.

- Click **Delete** ( ) to delete a rule.
- Click a rule and drag it to a new location to move it. You cannot drag rules between the interface and global lists. Instead you must edit the rule to change the interface/global setting. The first rule in the list that matches a connection is applied to the connection.

Step 5 Click **OK** when you are finished editing the policy.

Step 6 Click **Save** on the **Advanced** window. The changes are not saved until you click **Save**.

Configure a service policy rule

Use this procedure to apply specific services and connection settings to defined traffic classes, enabling you to control and customize how network traffic is handled.

Service policy rules allow you to control and customize how network traffic is handled by applying specific services and connection settings to defined traffic classes.

Before you begin

Go to **Objects** > **Object Management** > **Access List** > **Extended** and create an extended access list that defines the traffic to which the rule applies. The rule is applied to any connections that match Allow rules in the extended access list. Define the ACL rules precisely, so that your service policy rule applies to only the traffic that requires the service.

If you are creating an interface-based rule, you must also have configured the interfaces on the assigned devices and added them to security zones or interface groups.

Follow these steps to configure a service policy rule:

Procedure

Step 1 If you are not already in the Threat Defense service policy dialog box, choose **Policies** > **Access Control heading** > **Access Control**, edit the access control policy, select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, then edit the **Threat Defense Service Policy**.

Step 2 Do any of the following:

- Click **Add Rule** to create a new rule.
- Click **Edit** ( ) to edit an existing rule.

The service policy rule wizard opens to step you through the process of configuring the rule.

Step 3 In the **Interface Object** step, select the option that defines the interfaces that will use the policy.

- **Apply Globally**—Select this option to create a global rule, which applies to all interfaces.
- **Select Interface Objects**—Select this option to create an interface-based rule. Then, select the security zones or interface objects that contain the desired interfaces, and click **Add** to move them to the selected list. The service policy rule will be configured on each interface contained in the selected objects; it is not configured on the zone/group itself.

Click **Next** when the interface criterion is complete.

Step 4 In the **Traffic Flow** step, select the extended ACL object that defines the connections to which the rule applies, then click **Next**.

Step 5 In the **Connection Setting** step, configure the services to apply to this traffic class.

- **Enable TCP State Bypass** (TCP connections only)—Implement TCP state bypass. Connections subject to TCP state bypass are not inspected by any inspection engines, and they bypass all TCP state checking and TCP normalization. For detailed information, see Bypass TCP state checks for asymmetrical routing (TCP state bypass), on page 2714.

  Note
  Use TCP state bypass for troubleshooting purposes or when asymmetric routing cannot be resolved. This feature disables multiple security features, which can cause a high number of connections if you do not implement it properly with a narrowly-defined traffic class.

- **Randomize TCP Sequence Number** (TCP connections only)—Whether to enable or disable TCP sequence number randomization. Randomization is enabled by default. For more information, see Disable TCP sequence randomization, on page 2718.

- **Enable Decrement TTL** (TCP connections only)—Decrement the time-to-live (TTL) on packets that match the class. If you decrement time to live, packets with a TTL of 1 will be dropped, but a connection will be opened for the session on the assumption that the connection might contain packets with a greater TTL. Note that some packets, such as Open Shortest Path First (OSPF) hello packets, are sent with TTL = 1, so decrementing time to live can have unexpected consequences.

  Note
  If you want the Firewall Threat Defense device to appear on traceroutes, you must configure the decrement TTL option and also set the Internet Control Message Protocol (ICMP) unreachable rate limit in the platform settings policy. See Configure the Firewall Threat Defense device to appear on traceroutes, on page 2723.

- **Connections**—Limits for the number of connections allowed for the entire class. You can configure these options:

  - **Maximum TCP and UDP** (TCP/UDP connections only)—The maximum number of simultaneous connections that are allowed, between 0 and 2000000, for the entire class. For TCP, this count applies to established connections only. The default is 0, which allows unlimited connections. Because the limit is applied to a class, one attacking host can consume all the connections and leave none for the rest of the hosts that are matched to the class. Set the per-client limit to ameliorate this problem.

  - **Maximum Embryonic** (TCP connections only)—The maximum number of simultaneous embryonic TCP connections (those that have not finished the TCP handshake) allowed, between 0 and 2000000. The default is 0, which allows unlimited connections. By setting a non-zero limit, you enable TCP intercept, which protects inside systems from a denial of service (DoS) attack perpetrated by flooding an interface with TCP SYN packets. Also set the per-client options to protect against SYN flooding. For more information, see Protect servers from a SYN flood DoS attack (TCP intercept), on page 2720.

- **Connections Per Client**—Limits for the number of connections allowed for a given client (source IP address). You can configure these options:

- **Maximum TCP and UDP** (TCP/UDP connections only)—The maximum number of simultaneous connections allowed per client, between 0 and 2000000. For TCP, this includes established, half-open (embryonic), and half-closed connections. The default is 0, which allows unlimited connections. This option restricts the maximum number of simultaneous connections that are allowed for each host that is matched to the class.
- **Maximum Embryonic** (TCP connections only)—The maximum number of simultaneous embryonic TCP connections allowed per client, between 0 and 2000000. The default is 0, which allows unlimited connections. For more information, see Protect servers from a SYN flood DoS attack (TCP intercept), on page 2720.

- **Connections SYN Cookie MSS**—The server maximum segment size (MSS) for SYN-cookie generation for embryonic connections upon reaching the embryonic connections limit, from 48 to 65535 . The default is 1380. This setting is meaningful only if you configure **Maximum Embryonic** for connections or per-client, or both.
- **Connections Timeout**—The timeout settings to apply to the traffic class. These timeouts override the global timeouts defined in the platform settings policy. You can configure these options:
  - **Embryonic** (TCP connections only)—The timeout period until a TCP embryonic (half-open) connection is closed, between 0:0:5 and 1193:00:00. The default is 0:0:30.
  - **Half Closed** (TCP connections only)—The idle timeout period until a half-closed connection is closed, between 0:0:30 and 1193:0:0. The default is 0:10:0. Half-closed connections are not affected by dead connection detection (DCD). Also, the system does not send a reset when taking down half-closed connections.
  - **Idle** (TCP, UDP, ICMP, IP connections)—The idle timeout period after which an established connection of any protocol closes, between 0:0:1 and 1193:0:0. The default is 1:0:0, unless you select the TCP state bypass option, where the default is 0:2:0.
  - **Reset Connection Upon Timeout** (TCP connections only)—Whether to send a TCP RST packet to both end systems after idle connections are removed.
- **Detect Dead Connections** (TCP connections only)—Whether to enable dead connection detection (DCD). Before expiring an idle connection, the system probes the end hosts to determine if the connection is valid. If both hosts respond, the connection is preserved, otherwise the connection is freed. When operating in transparent firewall mode, you must configure static routes for the endpoints. You cannot configure DCD on connections that are also offloaded, so do not configure DCD on connections you are fast-pathing in the prefilter policy. Use the **show conn detail** command in the Firewall Threat Defense CLI to track how many DCD probes have been sent by the initiator and responder.

  Configure these options:
  - **Detection Timeout**—The time duration in hh:mm:ss format to wait after each unresponsive DCD probe before sending another probe, between 0:0:1 and 24:0:0. The default is 0:0:15.

    For systems that are operating in a cluster or high-availability configuration, we recommend that you do not set the interval to less than one minute (0:1:0). If the connection needs to be moved between systems, the changes required take longer than 30 seconds, and the connection might be deleted before the change is accomplished.
  - **Detection Retries**—The number of consecutive failed retries for DCD before declaring the connection dead, from 1 to 255. The default is 5.

Step 6 Click **Finish** to save your changes.

The rule is added to the bottom of the appropriate list, either Interfaces or Global. Global rules are matched in top-down order. Rules in the Interfaces list are matched in top down order for each interface object. Place rules for narrowly-defined traffic classes above broader rules, to ensure the right services get applied. You can move rules within each list by using drag and drop. You cannot move rules between lists.

Bypass TCP state checks for asymmetrical routing (TCP state bypass)

TCP state bypass allows traffic to bypass stateful inspection checks on Firewall Threat Defense devices.

Implementation considerations

If you have an asymmetrical routing environment in your network, where the outbound and inbound flow for a given connection can go through two different Firewall Threat Defense devices, you need to implement TCP state bypass on the affected traffic.

However, TCP state bypass weakens the security of your network, so you should apply bypass on very specific, limited traffic classes.

The asymmetrical routing problem

The asymmetrical routing problem occurs when return traffic for a connection goes through a different firewall than the outbound traffic. This causes packets to be dropped because they do not match existing connection entries in the fast path table.

Traffic inspection and fast path operation

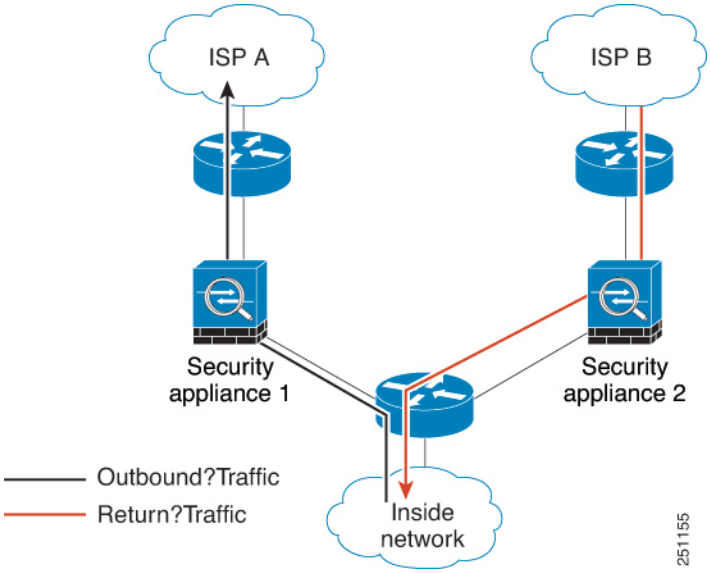
By default, the Firewall Threat Defense inspects all traffic using the Adaptive Security Algorithm. The Firewall Threat Defense allows or drops traffic based on the security policy. The Firewall Threat Defense maximizes the firewall performance by checking the state of each packet (new connection or established connection). It assigns new connection SYN packets to the session management path. It assigns established connections to the fast path. It assigns packets requiring advanced inspection to the control plane path.

The Firewall Threat Defense allows TCP packets that match existing connections to pass through via the fast path without rechecking every aspect of the security policy. This feature improves performance. However, the fast path establishment method can interfere with asymmetric routing solutions. The method uses SYN packets and performs checks such as TCP sequence number validation. Both outbound and inbound flows of a connection must pass through the same Firewall Threat Defense device.

Asymmetric routing scenario

For example, a new connection goes to Security Appliance 1. The SYN packet goes through the session management path, and an entry for the connection is added to the fast path table. If subsequent packets of this connection go through Security Appliance 1, then the packets match the entry in the fast path, and are passed through. But if subsequent packets go to Security Appliance 2, where no SYN packet went through the session management path, then no entry exists in the fast path for the connection, and Security Appliance 2 drops the packets. The following figure shows an asymmetric routing example where the outbound traffic goes through a different Firewall Threat Defense than the inbound traffic:

Figure 392: Asymmetric routing



If you have asymmetric routing configured on upstream routers and traffic alternates between two Firewall Threat Defense devices, then you can configure TCP state bypass for specific traffic. TCP state bypass alters the way sessions are established in the fast path and disables the fast path checks. This feature treats TCP traffic similarly to a UDP connection: when a non-SYN packet matching the specified networks enters the Firewall Threat Defense device, and there is not a fast path entry, then the packet goes through the session management path to establish the connection in the fast path. Once in the fast path, the traffic bypasses the fast path checks.

Guidelines and limitations for TCP state bypass

TCP state bypass unsupported features

The system does not support these features when you use TCP state bypass:

- Application inspection—Inspection requires both inbound and outbound traffic to go through the same Firewall Threat Defense, so inspection is not applied to TCP state bypass traffic.
- Snort inspection—Inspection requires both inbound and outbound traffic to go through the same device. However, Snort inspection is not automatically bypassed for TCP state bypass traffic. You must also configure a prefilter fastpath rule for the same traffic class for which you configure TCP state bypass. Otherwise, packets might be dropped unexpectedly because the TCP Normalizer is also not engaged.
- TCP Intercept, maximum embryonic connection limit, TCP sequence number randomization—The Firewall Threat Defense does not keep track of the state of the connection, so these features are not applied.
- TCP normalization—The TCP normalizer is disabled.
- TLS server identity discovery cannot be used with TCP state bypass on an inline or inline tap interface.

TCP state bypass NAT guidelines

Because the translation session is established separately for each device, be sure to configure static NAT on both devices for TCP state bypass traffic. If you use dynamic NAT, the address chosen for the session on Device 1 will differ from the address chosen for the session on Device 2.

Configure TCP state bypass

Configure TCP state bypass to handle asymmetrical routing environments where standard TCP state checking would interfere with traffic flow.

To bypass TCP state checking in asymetrical routing environments, carefully define a traffic class that applies only to the affected hosts or networks. Then, enable TCP state bypass on the traffic class using a service policy. You must also configure a corresponding prefilter fastpath policy for the same traffic to ensure that the traffic bypasses inspection.

Because bypass reduces the security of the network, limit its application as much as possible.

Before you begin

Follow these steps to configure TCP state bypass:

Procedure

Step 1 Create the extended ACL that defines the traffic class.

For example, to define a traffic class for TCP traffic from 10.1.1.1 to 10.2.2.2, do the following:

a) Choose **Objects** > **Object Management** > **Access List** > **Extended**.
b) Click **Add Extended Access List**.
c) Enter a **Name** for the object, for example, bypass.
d) Click **Add** to add a rule.
e) Keep **Allow** for the action.
f) Enter 10.1.1.1 beneath the **Source** list and click **Add**. Then, enter 10.2.2.2 beneath the **Destination** list, and click **Add**.
g) Click **Port**, select **TCP (6)** beneath the **Selected Source Ports** list, then click **Add**. Do not enter a port number. Add TCP as the protocol. This will cover all ports.
h) In the Extended Access List Entry dialog box, click **Add** to add the rule to the ACL.
i) In the Extended Access List Object dialog box, click **Save** to save the ACL object.

Step 2 Configure the TCP state bypass service policy rule.

For example, to configure TCP state bypass for this traffic class globally, do the following:

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy assigned to the devices that require this service.
b) Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, and click **Edit** (✎) for the **Threat Defense Service Policy**.
c) Click **Add Rule**.
d) Select **Global** and click **Next**.
e) Select the extended ACL object you created for this rule and click **Next**.
f) Select **Enable TCP State Bypass**.
g) (Optional.) Adjust the **Idle** timeout for bypassed connections. The default is 2 minutes.

h) Click **Finish** to add the rule. If necessary, drag and drop the rule to the desired position in the service policy.

Click **OK** to save the changes to the service policy.

Click **Save** on **Advanced** to save the changes to the access control policy.

Step 3 Configure a prefilter fastpath rule for the traffic class.

You cannot use the ACL object in the prefilter rule, so you need to recreate the traffic class either directly in the prefilter rule, or by first creating network objects that define the class.

This procedure assumes that you already have a prefilter policy attached to the access control policy. If you have not created a prefilter policy yet, go to **Policies** > **Prefilter** and first create the policy. You can then follow this procedure to attach it to the access control policy and create the rule.

Keeping with our example, this procedure creates a fastpath rule for TCP traffic from 10.1.1.1 to 10.2.2.2.

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy that has the TCP bypass service policy rule.
b) Click the link for the **Prefilter Policy**, which is to the left immediately under the policy description.
c) In the Prefilter Policy dialog box, select the policy to assign to the device if the correct one is not already selected. Do not click OK.

Because you cannot add rules to the default prefilter policy, you must choose a custom policy.

d) In the Prefilter Policy dialog box, click **Edit** ( ). This action opens a new browser window where you can edit the policy.
e) Click **Add Prefilter Rule** and configure a rule with these properties.

- **Name**—Any name that you find meaningful will do, such as TCPBypass.
- **Action**—Select **Fastpath**.
- **Interface Objects**—If you configured TCP state bypass as a global rule, leave the default, any, for both source and destination. If you created an interface-based rule, select the same interface objects you used for rule in the **Source Interface Objects** list, and keep any as the destination.
- **Networks**—Add 10.1.1.1 to the **Source Networks** list, and 10.2.2.2 to the **Destination Networks** list. You can either use network objects or manually add the addresses.
- **Ports**—Under **Selected Source Ports**, select TCP(6), **do not enter a port**, and click **Add**. This applies the rule to all TCP traffic, regardless of TCP port number.

f) Click **Add** to add the rule to the prefilter policy.
g) Click **Save** to save your changes to the prefilter policy.

You can now close the prefilter edit window and return to the access control policy edit window.

h) In the access control policy edit window, the Prefilter Policy dialog box should still be open. Click **OK** to save your changes to the prefilter policy assignment.
i) Click **Save** on the access control policy to save the changed prefilter policy assignment, if you changed it.

You can now deploy the changes to the affected devices.

Disable TCP sequence randomization

Disable TCP initial sequence number randomization to prevent issues with data scrambling and to support specific network configurations that require non-randomized sequence numbers.

Each TCP connection has two initial sequence numbers (ISN): one generated by the client and one generated by the server. The Firewall Threat Defense device randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session. However, TCP sequence randomization effectively breaks TCP SACK (Selective Acknowledgement), as the sequence numbers the client sees are different from what the server sees.

You can disable TCP initial sequence number randomization if data is getting scrambled or if there is another necessity. This is useful in these situations:

- If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
- If you use eBGP multi-hop through the device, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- If you use a WAAS device that requires the device not to randomize the sequence numbers of connections.
- If you enable hardware bypass for the ISA 3000, and TCP connections are dropped when the ISA 3000 is no longer part of the data path.

Procedure

Step 1 Create the extended ACL that defines the traffic class.

For example, to define a traffic class for TCP traffic from any host to 10.2.2.2, do the following:

a) Choose **Objects** > **Object Management** > **Access List** > **Extended**.
b) Click **Add Extended Access List**.
c) Enter a **Name** for the object, for example, preserve-sq-no.
d) Click **Add** to add a rule.
e) Keep **Allow** for the action.
f) Leave the **Source** list empty, enter 10.2.2.2 beneath the **Destination** list, and click **Add**.
g) Click **Port**, select **TCP (6)** beneath the **Selected Source Ports** list, and click **Add**. Do not enter a port number, simply add TCP as the protocol, which will cover all ports.
h) Click **Add** on the Extended Access List Entry dialog box to add the rule to the ACL.
i) Click **Save** on the Extended Access List Object dialog box to save the ACL object.

Step 2 Configure the service policy rule that disables TCP sequence number randomization.

For example, to disable randomization for this traffic class globally, do the following:

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy assigned to the devices that require this service.
b) Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, and click **Edit** (✎) for the **Threat Defense Service Policy**.

c) Click **Add Rule**.
d) Select **Global** and click **Next**.
e) Select the extended ACL object you created for this rule and click **Next**.
f) Deselect **Randomize TCP Sequence Number**.
g) (Optional.) Adjust the other connection options as needed.
h) Click **Finish** to add the rule. If necessary, drag and drop the rule to the desired position in the service policy.

Click **OK** to save the changes to the service policy.

Click **Save** on **Advanced** to save the changes to the access control policy.

You can now deploy the changes to the affected devices.

Monitor service policies

You can monitor service-policy related information using the device CLI. Use these commands to get the information you need.

Commands for monitoring service policies.

- **show conn** [**detail**]

  Shows connection information. Detailed information uses flags to indicate special connection characteristics. For example, the "b" flag indicates traffic subject to TCP state bypass.

  When you use the **detail** keyword, you can see information about dead connection detection (DCD) probing, which shows how often the connection was probed by the initiator and responder. For example, the connection details for a DCD-enabled connection would look like this:

```
TCP dmz: 10.5.4.11/5555 inside: 10.5.4.10/40299,
    flags UO , idle 1s, uptime 32m10s, timeout 1m0s, bytes 11828,
cluster sent/rcvd bytes 0/0, owners (0,255)
  Traffic received at interface dmz
        Locally received: 0 (0 byte/s)
  Traffic received at interface inside
        Locally received: 11828 (6 byte/s)
  Initiator: 10.5.4.10, Responder: 10.5.4.11
  DCD probes sent: Initiator 5, Responder 5
```

- **show service-policy**

  Shows service policy statistics, including dead connection detection (DCD) statistics.

- **show threat-detection statistics top TCP-intercept** [**all** | **detail**]

  View the top 10 protected servers under attack. The **all** keyword shows the history data of all the traced servers. The **detail** keyword shows history sampling data. The system samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

Examples for service policy rules

Provides examples of service policy rules to help users understand how to configure and implement service policies effectively.

Protect servers from a SYN flood DoS attack (TCP intercept)

This task configures TCP intercept to protect servers from SYN flood denial of service attacks by limiting embryonic connections and enabling monitoring statistics.

A SYN-flooding denial of service (DoS) attack occurs when an attacker sends a series of SYN packets to a host. These packets usually originate from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests from legitimate users.

You can limit the number of embryonic connections to help prevent SYN flooding attacks. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination.

When the embryonic connection threshold of a connection is crossed, the Firewall Threat Defense acts as a proxy for the server. It generates a SYN-ACK response to the client SYN request using the SYN cookie method. This approach prevents the connection from being added to the SYN queue of the targeted host. The SYN cookie is the initial sequence number returned in the SYN-ACK that is constructed from MSS, time stamp, and a mathematical hash of other items to essentially create a secret. If the Firewall Threat Defense receives an ACK from the client with the correct sequence number, and within the valid time window, it authenticates the client and allows the connection to the server. The component that performs the proxy is called TCP intercept.

Setting connection limits can protect a server from a SYN flood attack. You can optionally enable TCP intercept statistics and monitor the results of your policy. The following procedure explains the end-to-end process.

Before you begin

- Ensure that you set the embryonic connection limit lower than the TCP SYN backlog queue on the server that you want to protect. Otherwise, valid clients can no longer access the server during a SYN attack. To determine reasonable values for embryonic limits, carefully analyze the capacity of the server, the network, and server usage.
- Depending on the number of CPU cores on your Secure Firewall Threat Defense device model, the maximum concurrent and embryonic connections can exceed the configured numbers due to the way each core manages connections. In the worst case scenario, the device allows up to n-1 extra connections and embryonic connections, where n is the number of cores. For example, if your model has 4 cores, if you configure 6 concurrent connections and 4 embryonic connections, you could have an additional 3 of each type. To determine the number of cores for your model, enter the **show CPU core** command in the device CLI.

Follow these steps to protect servers from a SYN flood DoS attack using TCP intercept:

Procedure

Step 1 Create the extended ACL that defines the traffic class, which is the list of servers you want to protect.

For example, to define a traffic class to protect the web servers with the IP addresses 10.1.1.5 and 10.1.1.6:

a) Choose **Objects** > **Object Management** > **Access List** > **Extended**.
b) Click **Add Extended Access List**.
c) Enter a **Name** for the object, for example, protected-servers.
d) Click **Add** to add a rule.
e) Keep **Allow** for the action.
f) Leave the **Source** list empty, enter 10.1.1.5 beneath the **Destination** list, and click **Add**.
g) Also enter 10.1.1.6 beneath the **Destination** list and click **Add**.
h) Click **Port**, select **HTTP** in the available ports list, and click **Add to Destination**. If your server also support HTTPS connections, also add that port.
i) Click **Add** to add the rule to the ACL.

Click **Save** to save the ACL object.

Step 2 Configure the service policy rule that sets embryonic connection limits.

For example, to set the total concurrent embryonic limit to 1000 connections, and the per-client limit to 50 connections, do the following:

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy assigned to the devices that require this service.
b) Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, and click **Edit** ( ) for the **Threat Defense Service Policy**.
c) Click **Add Rule**.
d) Select **Global** and click **Next**.
e) Select the extended ACL object you created for this rule and click **Next**.
f) Enter 1000 for **Connections** > **Maximum Embryonic**.
g) Enter 50 for **Connections Per Client** > **Maximum Embryonic**.

(Optional.) Adjust the other connection options as needed.

h) Click **Finish** to add the rule. If necessary, drag and drop the rule to the desired position in the service policy.

Click **OK** to save the changes to the service policy.

Click **Save** on **Advanced** to save the changes to the access control policy.

Step 3 (Optional.) Configure the rates for TCP intercept statistics.

TCP intercept uses these options to determine the rate for collecting statistics. All options have default values, so if these rates suit your needs, you can skip this step.

- Rate Interval—The size of the history monitoring window, between 1 and 1440 minutes. The default is 30 minutes. During this interval, the system samples the number of attacks 30 times.
- Burst Rate—The threshold for syslog message generation, between 25 and 2147483647. The default is 400 per second. When the burst rate is exceeded, the device generates syslog message 733104.
- Average Rate—The average rate threshold for syslog message generation, between 25 and 2147483647. The default is 200 per second. When the average rate is exceeded, the device generates syslog message 733105.

If you want to adjust these options, do this:

a) Choose **Objects** > **Object Management** > **FlexConfig** > **Text Object**.
b) Click **Edit** ( ) for the **threat_defense_statistics** system-defined object.
c) Although you can directly change the values, the recommended approach is to open the **Override** section and click **Add** to create a device override.
d) Select the devices to which you will assign the service policy (through the access control policy assignment) and click **Add** to move them to the selected list.
e) Click **Override**.
f) The object must have 3 entries, so click **Count** as needed until you get 3.
g) Enter the values you need, in order from 1-3, as rate interval, burst rate, and average rate. Consult the object description to verify you enter the values in the right order.
h) Click **Add** in the Object Override dialog box.
i) Click **Save** in the Edit Text Object dialog box.

Step 4 Enable TCP intercept statistics.

You must configure a FlexConfig policy to enable TCP intercept statistics.

a) Choose **Devices** > **FlexConfig**.
b) If you already have a policy assigned to the devices, edit it. Otherwise, create a new policy and assign it to the affected devices.
c) Select the **Threat_Detection_Configure** object in the **Available FlexConfig** list and click >>. The object is added to the **Selected Append FlexConfigs** list.
d) Click **Save**.
e) (Optional.) You can verify that you have the right settings by clicking **Preview Config** and selecting one of the devices.

The system generates the CLI commands that will be written to the device during the next deployment. These commands will include those needed for the service policy as well as those needed for threat detection statistics. Scroll to the bottom of the preview to see the appended CLI. The TCP intercept statistics command should look something like this, if you use the default values (line break added for clarity):

```
###Flex-config Appended CLI ###

threat-detection statistics tcp-intercept rate-interval 30
       burst-rate 400 average-rate 200
```

Step 5 You can now deploy the changes to the affected devices.

Step 6 Monitor the TCP intercept statistics from the device CLI using these commands:

- **show threat-detection statistics top TCP-intercept** [**all** | **detail**]—View the top 10 protected servers under attack. The **all** keyword shows the history data of all the traced servers. The **detail** keyword shows history sampling data. The system samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

  Note
  You can use the **shun** command to block attacking host IP addresses. To remove the block, use the **no shun** command.

- **clear threat-detection statistics TCP-intercept**—Erases TCP intercept statistics.

Example:

```
hostname(config)# show threat-detection statistics top tcp-intercept
Top 10 protected servers under attack (sorted by average rate)
Monitoring window size: 30 mins    Sampling interval: 30 secs
<Rank> <Server IP:Port> <Interface> <Ave Rate> <Cur Rate> <Total> <Source IP (Last Attack
Time)>
--------------------------------------------------------------------------------
1    10.1.1.5:80 inside 1249 9503 2249245 <various> Last: 10.0.0.3 (0 secs ago)
2    10.1.1.6:80 inside 10 10 6080 10.0.0.200 (0 secs ago)
```

Configure the Firewall Threat Defense device to appear on traceroutes

Make the Firewall Threat Defense device appear as a hop in traceroute results by decrementing the time-to-live on packets and adjusting ICMP rate limits.

By default, the Firewall Threat Defense device does not appear on traceroutes as a hop. To make it appear, you need to decrement the time-to-live on packets that pass through the device, and increase the rate limit on ICMP unreachable messages. To accomplish this, you must configure a service policy rule and adjust the ICMP platform settings policy.

Note If you decrement time to live, packets with a TTL of 1 will be dropped, but a connection will be opened for the session on the assumption that the connection might contain packets with a greater TTL. Note that some packets, such as OSPF hello packets, are sent with TTL = 1, so decrementing time to live can have unexpected consequences. Keep these considerations in mind when defining your traffic class.

Procedure

Step 1 Create the extended ACL that defines the traffic class for which to enable traceroute reporting.

For example, to define a traffic class for all addresses, but excluding OSPF traffic, do the following:

a) Choose **Objects** > **Object Management** > **Access List** > **Extended**.
b) Click **Add Extended Access List**.
c) Enter a **Name** for the object, for example, traceroute-enabled.
d) Click **Add** to add a rule to exclude OSPF.
e) Change the action to **Block**, click **Port**, select **OSPFIGP (89)** as the protocol beneath the **Destination Ports** list, and click **Add** to add the protocol to the selected list.
f) Click **Add** on the Extended Access List Entry dialog box to add the OSPF rule to the ACL.
g) Click **Add** to add a rule to include all other connections.
h) Keep **Allow** for the action, and leave both the Source and Destination lists empty.
i) Click **Add** on the Extended Access List Entry dialog box to add the rule to the ACL.

Ensure that the OSPF deny rule is above the Allow Any rule. Drag and drop to move the rules if necessary.

Click **Save** on the Extended Access List Object dialog box to save the ACL object.

Step 2 Configure the service policy rule that decrements the time-to-live value.

For example, to decrement time-to-live globally, do the following:

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy assigned to the devices that require this service.
b) Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, and click **Edit** (✎) for the **Threat Defense Service Policy**.
c) Click **Add Rule**.
d) Select **Global** and click **Next**.
e) Select the extended ACL object you created for this rule and click **Next**.
f) Select **Enable Decrement TTL**.
g) (Optional.) Adjust the other connection options as needed.
h) Click **Finish** to add the rule. If necessary, drag and drop the rule to the desired position in the service policy.

   Click **OK** to save the changes to the service policy.

   Click **Save** on **Advanced** to save the changes to the access control policy.

You can now deploy the changes to the affected devices.

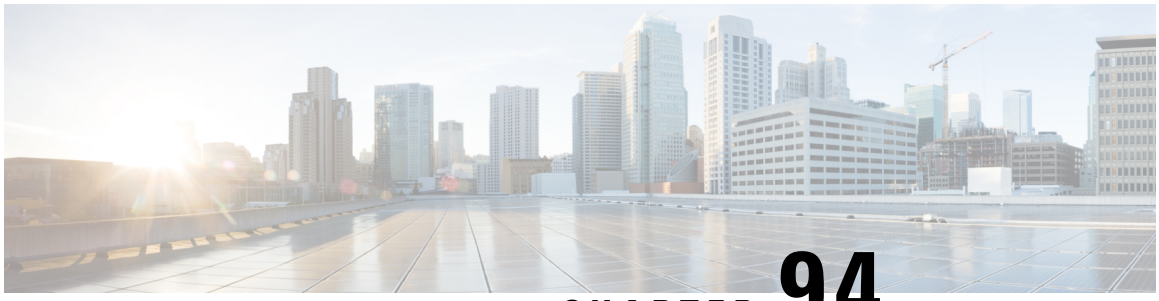
Step 3 Increase the rate limit on ICMP unreachable messages.

a) Choose **Devices** > **Platform Settings**.
b) If you already have a policy assigned to the devices, edit it. Otherwise, create a new Threat Defense platform settings policy and assign it to the affected devices.
c) Select **ICMP** from the table of contents.
d) Increase the **Rate Limit**, for example, to 50. You might also want to increase the **Burst Size**, for example, to 10, to ensure enough responses are generated within the rate limit.

   You can leave the ICMP rules table empty, it is not related to this task.

e) Click **Save**.

Step 4 You can now deploy the changes to the affected devices.



CHAPTER 94

Threat Detection

Threat detection's portscan detector is a mechanism designed to help you detect and prevent portscan activity in all types of traffic to protect networks from eventual attacks. Portscan traffic can be detected efficiently in both allowed and denied traffic.

Portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker determines the types of network protocols or services a host supports and sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

- Portscan detection and prevention, on page 2725
- Requirements and prerequisites for threat detection, on page 2727
- Limitations for threat detection configuration, on page 2728
- Best practices for portscan prevention, on page 2729
- Configure portscan detection and prevention, on page 2729
- Monitoring threat detection, on page 2731

Portscan detection and prevention

Portscan detection and prevention is a threat detection system that:

- identifies port scan activity and issues events when they are found,
- optionally prevents port scans by automatically blocking scanners, and
- sends you events and also blocks the attacker for a duration that you set.

Pre-defined sensitivity levels for portscan detection

When configuring detection settings, you select from these pre-defined sensitivity levels. Except for custom, each level has pre-set values for each protocol for the number of ports (TCP, UDP), protocols (IP), or hosts (TCP, UDP, IP, ICMP) that must be scanned within a set time interval (expressed in seconds). Also, all types of scans and sweeps are enabled.

Note When counting ports or protocols, threat detection increments the number if the port or protocol in the current packet differs from the previous packet. For example, if you have an application that opens connections in 10 set ports randomly, the total number of ports scanned could mount so quickly that your port number will be exceeded within the interval. The system does not count only unique ports.

Exceeding the number within the interval can indicate a scanning attack. Portscan events are generated only when the port/protocol/host numbers are exceeded for the moving time interval window.

These sensitivity levels determine portscan detection behavior:

- **Low**—This level uses the shortest time window for portscan detection, coupled with high counts for ports, protocols, and hosts. Thus, you should see portscan events for the most aggressive scanners only. Select this sensitivity level to suppress false positives, but remember that some types of port scans, such as slow or filtered scans, might be missed. For more detail on how low sensitivity detection works, see Detection in the low sensitivity level, on page 2727.
- **Medium**—This level uses moderate values for both the interval and port/protocol/host counts. However, very active hosts such as network address translators and proxies might generate false positives. Add such hosts to the ignore scanner list. This is the default sensitivity level and a good place to start.
- **High**—This level uses a much longer time window for portscan detection, coupled with lower counts for ports, protocols, and hosts. With this level, you are most likely to see events for even the least aggressive port scans or sweeps, so you are more likely to notice all attackers. However, this level would likely result in the most portscan events issued, and potentially the highest number of false positives.
- **Custom**—If you want to configure any setting differently than one of the predefined sensitivity levels, or disable a particular type of scan or sweep, the level automatically switches to custom. If you want to adjust the options, first select the level that most closely matches what you want, then edit the values as appropriate.

Table 145: Sensitivity level settings

Setting	Low	Medium	High
Interval (TCP/UDP/IP/ICMP)	60 seconds	90 seconds	600 seconds (10 minutes)
TCP/UDP portscan - Number of Ports	120 seconds	90 seconds	60 seconds
TCP/UDP portsweep - Number of Hosts	180 seconds	150 seconds	100 seconds
IP protocol scan - Number of Protocols	30 seconds	15 seconds	10 seconds
IP protocol sweep - Number of Hosts	25 seconds	20 seconds	10 seconds
ICMP host sweep - Number of Hosts	50 seconds	30 seconds	20 seconds

Detection in the low sensitivity level

When you select the low sensitivity level, the system tracks negative responses for TCP, UDP, and ICMP initial packets. An alert is triggered only when the number of unsuccessful connections exceeds the rejection threshold (10% for low sensitivity) and the port or IP protocol count exceeds the configured threshold. This mitigates false positives.

Low sensitivity level detection behavior

Rejection threshold applies to low sensitivity (or equivalent custom settings) only; it does not apply to other sensitivity levels or their custom equivalents.

When there is a mix of allowed and blocked traffic, the number of rejected ports or hosts is calculated based on the difference between allowed and blocked traffic. In the case of only blocked traffic, the rejection threshold is not considered.

These criteria are not used for UDP/ICMP connections on interfaces configured in inline sets.

Low sensitivity level detection scenarios

In low sensitivity mode, the port count threshold is 120. Thus, the rejection count threshold is 10% of 120, which is 12. Here are examples that show how the system issues portscan events with this configuration:

- An attacker initiates connections with 131 ports of the target and the target positively acknowledges all the initiations. Port count = 131, which is greater than the threshold, but a portscan alert is not triggered because there are no negative acknowledgements.
- An attacker initiates connections with 131 ports of the target and the target positively acknowledges 121 initiations and negatively acknowledges 10 initiations. Port count = 131, which is greater than the threshold, but reject port count = 10, which is lesser than the rejection threshold. Therefore, a portscan alert is not triggered.
- An attacker initiates connections with 134 ports of the target and the target positively acknowledges 121 initiations and negatively acknowledges 13 initiations. Port count = 134, which is greater than the threshold, and reject port count = 13 is also higher than the rejection threshold. Therefore a portscan alert is triggered.

Requirements and prerequisites for threat detection

Model support

Firewall Threat Defense running version 7.2+ and Snort 3.

Supported domains

Any

User roles

Admin

Access Admin

Network Admin

Limitations for threat detection configuration

Threat detection requires Snort 3 and has specific version requirements and configuration interactions that vary by management platform.

- Threat detection requires Snort 3. The Network Analysis Policy (NAP) port scan configuration is always ignored for devices running Snort 3; you must configure port scan using threat detection. For Snort 2, you can configure port scan through the NAP policy only. If there are Snort 2 devices assigned to the access control policy, the threat detection settings will not be deployed to those unsupported devices.
- Threat detection requires Snort 3. The managed device must be running version 7.2 or higher. For Snort 2, or devices running versions lower than 7.2, you can configure port scan through the NAP policy. Note that the threat detection feature is not the same as the port scan feature in the NAP policy. If there are non-Snort 3 devices running version 7.2 or higher that are assigned to the access control policy, the threat detection settings are not deployed to those unsupported devices.
- If you configure port scan in the NAP policy on a device running 7.1 or lower, that configuration is not translated to the threat detection feature on upgrade to 7.2. You must manually configure threat detection. Although the NAP and threat detection port scan options are similar, they do not match one-to-one.
- If you configure threat detection, any port scan configuration in the NAP policy is ignored and not configured on the devices that support threat detection.
- The NAP port scan feature for Snort 3 is always ignored for devices running version 7.2 or higher. To configure port scanning, you must use the threat defense settings.
- Threat detection works on traffic that passes through the device only. It does not work on traffic directed to the device.
- In a high availability setup, port scanning statistics are not synchronized to the standby unit. However, blocked hosts are synchronized and continue to be blocked until the duration period expires in case of a failover.
- (Devices running Threat Defense versions 7.2 through 7.7). For nodes in a cluster, detection and prevention happen on the individual cluster node. That is, if node B detects and blocks traffic from a host, node A will not be aware of that action because port scan statistics are not synchronized across cluster nodes.
- (Devices running Threat Defense version 10.0+). For nodes in a cluster, detection and prevention happen at the cluster level. Port scans can be detected when they happen across nodes or in an individual node. Shunned hosts are shunned on all devices in the cluster. Shuns are released at the same time on all nodes. Statistics are available at the cluster level.
- For inline sets, or for interfaces that are configured to be part of an equal-cost multipath (ECMP) traffic zone, detection and prevention are done at the zone level. Port scan statistics for a host are accumulated across all interfaces of a zone. Similarly, when a host crosses configured thresholds, it is blocked across all interfaces of the corresponding zone.
- Although the port scan events generated by the threat detection feature are the same as the ones Snort issues for port scan, you do not need to enable port scanning intrusion rules to get the events. Threat detection works regardless of your intrusion policy implementation.

Best practices for portscan prevention

Portscan prevention mode can result in unintended traffic outage. In prevention mode, hosts are blocked from further scanning of networks on all protocols for the configured duration. Review the detection and prevention parameters carefully to ensure legitimate traffic is not blocked.

Before configuring portscan in prevention mode:

- Start using portscan in detection mode.
- Observe the generated portscan events.
- Tune the sensitivity level, monitored networks, ignore scanner list, and ignore target list. If a pre-defined sensitivity level does not work well for your situation, configure custom settings as needed.
- Repeat the process until false positives are eliminated and the event rate represents an accurate picture of port scanning in your network. Ensure that you are comfortable with blocking the remaining identified scanners.

Configure portscan detection and prevention

This task enables threat detection to watch for port scanning activity and optionally automatically block scanners for a period of time.

Portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker determines the types of network protocols or services a host supports and sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

You can enable threat detection to watch for port scanning activity and optionally, automatically block scanners for a period of time.

Before you begin

Fully Qualified Domain Name (FQDN), wildcard mask, any, any-ipv4, and any-ipv6 network objects are not supported for portscan configuration. These objects are not shown in the **Monitor**, **Ignore Scanner**, **Ignore Target**, and **Exclude** fields.

Procedure

Step 1 In the access control policy editor, click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line. Then, click **Edit** (✎) next to **Threat Detection**.

Step 2 In the **Threat Detection** window, select the **Portscan mode**:

- **Disable**—Turn off threat detection. This is the default mode. You can click **Revert to Defaults** to return to this unconfigured state.

- **Detection**—Perform portscan detection, but alert on problems only. Do not take action against potential attackers. We suggest you use this mode initially until you fine-tune the threat detection settings to avoid excessive false positives.
- **Prevention**—Perform portscan detection and actively block identified scanners, that is, hosts that are performing the port scan.

Step 3 Configure the **Traffic Selection** options.

The traffic selection options determine which networks are monitored, the type of connections monitored, and whether any scanners or target hosts should be exempted from the monitored networks. By default, the system monitors permitted connections on all networks.

- **Detection On Traffic**—Select the types of connection that will be monitored for portscan activity: **Permitted**, **Denied**, or **All** traffic. The default is **Permitted**.
- **Monitor**—Select the network objects that define the networks to monitor for portscan or sweep activity. The default is any network, IPv4 or IPv6. Use this option to limit scanning to untrusted networks.
- **Ignore Scanner**—Select the network objects that define the hosts or networks, from within the range of the monitored networks, that should be ignored. For example, if you have set up your own scanner to test your network, you can exempt the address of your scanner to avoid unnecessary reporting on the address. Do not include addresses that are outside the monitored networks, as these addresses are already ignored.
- **Ignore Target**—Select the network objects that define the hosts or networks that should be ignored as targets, that is, victims of a portscan or sweep.

Step 4 Click the **Configuration** tab and select the scanning sensitivity level.

The predefined sensitivity levels, **Low**, **Medium**, and **High**, set the port scanning options to values that are increasingly aggressive. For example, if you select low, you would expect to see fewer port scanning events, and you could potentially miss attackers more easily than if you selected medium or high. On the other hand, if you select high, you might see more events and also, potentially, more false positives. The default level is medium. For more information on these levels, see Pre-defined sensitivity levels for portscan detection, on page 2725.

As you select the levels, you can see the related values within the protocol sections: **TCP**, **UDP**, **IP**, and **ICMP**. If you change any of the preset values, or disable a type of scan, the sensitivity mode automatically changes to **Custom**.

Within each protocol section, the options are:

- **Interval**—The time range, in seconds, within which the configured values for portscan or portsweep are exceeded. For example, if you select 90 seconds, and 60 as the number of TCP portscan ports, a scanner would need to try 60 ports on a host within 90 seconds for it to be considered a portscan. The system generates events only if the number of ports, protocols, or hosts (for a portsweep) are exceeded within the specified interval.

  You can specify a range between 30 to 600 seconds. The longer the period, the more likely a host might be identified as a scanner.
- **Portscan** (TCP/UDP)—Select whether to monitor for port scanning against single hosts, and specify the number of ports that must be scanned within the interval to count as a portscan attack. The allowed range is 1-256.

- **Portsweep** (TCP/UDP)—Select whether to monitor for port sweeping against multiple hosts, and specify the number of hosts that must be scanned for a given port within the interval to count as a portsweep attack. The allowed range is 1-256.
- **Protocol Scan** (IP)—Select whether to monitor for protocol scanning against single hosts, and specify the number of protocols that must be scanned within the interval to count as a protocol scan attack. The allowed range is 1-255.
- **Protocol Sweep** (IP)—Select whether to monitor for protocol sweeping against multiple hosts, and specify the number of hosts that must be scanned for a given protocol within the interval to count as a protocol sweep attack. The allowed range is 1-256.
- **Hostsweep** (ICMP)—Select whether to monitor for ICMP host sweeping against multiple hosts, and specify the number of hosts that must be scanned within the interval to count as a hostsweep attack. The allowed range is 1-256.

Step 5 If you selected prevention mode, click the **Prevention** tab and configure the options.

In prevention mode, hosts are automatically blocked from further scanning of networks on all protocols for the configured duration. Review the detection and prevention parameters carefully to ensure legitimate traffic is not blocked.

- **Exclude**—Select the network objects that define the hosts or networks, from within the range of the monitored networks, that should be excluded from automatic blocking. Even if these hosts violate your scanning detection parameters, the system will not block them.
- **Duration**—How long, in seconds, automatically blocked scanner hosts should be prevented from sending traffic of any kind through the device. After the duration period ends, the hosts are automatically cleared and can again send traffic through the device. The allowed range is 600 to 2592000 seconds. The default is 3600 seconds (1 hour).

  If you need to manually unblock a host, SSH to the firewall that is blocking the host and use the **clear threat-detection portscan attacker** command.

Step 6 Click **OK** to save the threat detection settings.

Step 7 Click **Save** to save the access control policy.

What to do next

Deploy configuration changes.

Monitoring threat detection

The following topics explain how to monitor portscan activity.

View portscan alerts

Portscan activity is alerted through the existing portscan-specific intrusion events. Intrusion events with generator ID (GID) 122 and Snort ID from SIDs 1 through 27 are generated. For these events, the *(port_scan)*

string is prepended in the event messages. The events include packet information along with packet data containing the statistics that triggered the alert.

To viewportscan events, go to **Analysis** > **Unified Events**.

Portscan issues these events regardless of your intrusion policy or NAP configuration. Events are issued only when scanners exceed the number of configured ports, protocols, or hosts for the various types of scan or sweep within the configured time interval for the associated protocol. A port scan from one host generates one event per set interval as soon as the threshold is met. If the same host initiates a new port scan during the same interval, no event is reported.

Table 146: Portscan events

Portscan type	Intrusion event
TCP Regular, Decoy, Distributed Scan	122:1 (port_scan) TCP portscan
TCP Portsweep	122:3 (port_scan) TCP portsweep
TCP Distributed Scan	122:4 (port_scan) TCP distributed portscan
IP Regular, Decoy, Distributed Protocol Scan n	122:9 (port_scan) IP protocol scan
IP Protocol Sweep	122:11 (port_scan) IP protocol sweep
IP Distributed Scan	122:12 (port_scan) IP distributed protocol scan
UDP Regular, Decoy, Distributed Scan	122:17 (port_scan) UDP portscan
UDP Portsweep	122:19 (port_scan) UDP portsweep
UDP Distributed Scan	122:20 (port_scan) UDP distributed portscan
ICMP Sweep	122:25 (port_scan) ICMP sweep
Portscan Block	122:100 (port_scan) host blocked due to portscan activity

Monitor portscan on the firewall

To monitor portscan, log in to the device Command-Line Interface (CLI) and use these commands.

- **show threat-detection portscan** [**attacker** | **target** | **shun**]

  Shows the IP addresses of scanners, those that have been shunned (blocked), and hosts that have been targeted by scans or sweeps.

- **show threat-detection portscan statistics** [**host** [*ipv4_address* | *ipv6_address*]] [**protocol** {**tcp** | **udp** | **IP** | **icmp**}]

  Shows statistics related to the port scan system. You can specify host, protocol, or host and protocol to filter the output to the desired information.

- **clear threat-detection portscan** [**attacker** | **target** | **shun**] [*ipv4_address mask* | *ipv6_address/prefix*]

Manually unblocks scanners (attackers) or identified targets. Enter the command without parameters to clear all attackers, targets, or shunned hosts.

- **clear threat-detection portscan statistics** [**host** [*ipv4_address* | *ipv6_address*]] [**protocol** {**tcp** | **udp** | **IP** | **icmp**}]

  Erases statistics related to portscan, so that you can more clearly see the current state of scanning through this device. Enter the command without parameters to clear all statistics. Alternatively, specify a host, protocol, or host and protocol, to limit the reset to the specified items.

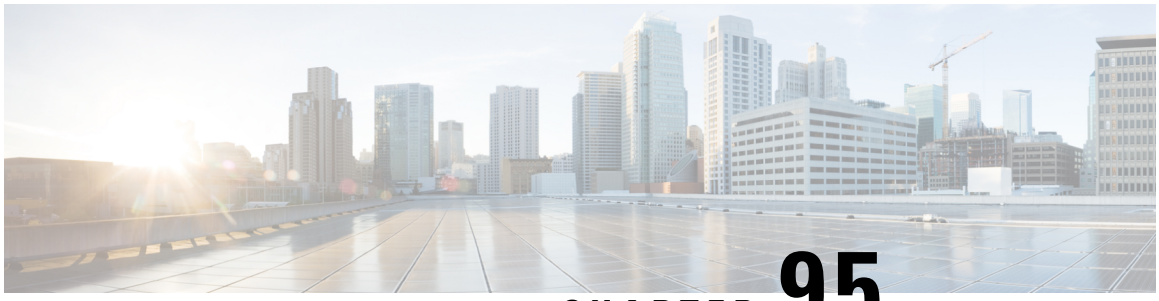
Unblocking a host

If you configure threat detection in prevention mode and the system blocks a host that you know is not an attacker, you can manually unblock the host before the host is automatically unblocked when the duration period expires.

To manually unblock the host, log in to the device CLI where the host is blocked and enter the clear threat-detection portscan attacker command. For example:

```
> clear threat-detection portscan attacker 10.2.0.100 255.255.255.255
1 tracker object deleted and 1 shun entry removed
```

Consider adding the host IP address to the exclude list in the prevention configuration.



CHAPTER 95

Elephant Flow Detection

Elephant flows are extremely large (in total bytes), continuous flows set up by a TCP (or other protocols) flow measured over a network link. By default, elephant flows are those larger than 1 GB/10 seconds. They can cause performance duress in Snort cores. Elephant flows are not numerous, but they can occupy a disproportionate share of the total bandwidth over a period of time. They can lead to problems, such as high CPU utilization, packet drops, and other issues.

From Cloud-Delivered Firewall Management Center 7.2.0 onwards (Snort 3 devices only), you can use the elephant flow feature to detect and remediate elephant flows, which helps to reduce system stress and resolve the mentioned issues.

- Elephant flow detection and remediation, on page 2735
- Elephant flow upgrade from Intelligent Application Bypass, on page 2736
- Configure elephant flow detection, on page 2736
- Use case - Configure Elephant Flow Detection Outcomes, on page 2739

Elephant flow detection and remediation

Elephant flow detection and remediation is a feature that

- detects elephant flows in network traffic,
- applies remediation actions to manage these flows, and
- requires Snort 3 as the detection engine to function.

Remediation actions

These remediation actions can be applied:

- **Bypass elephant flow**—You can configure elephant flow to bypass Snort inspection. If this is configured, Snort does not receive any packet from that flow.
- **Throttle elephant flow**—You can apply rate-limit to the flow and continue to inspect flows. The flow rate is calculated dynamically, and the system reduces the flow rate by 10 percent. Snort sends the verdict (with QoS flow reduced by 10 percent) to the firewall engine. If you choose to bypass all applications, including unidentified applications, you cannot configure the throttle action (rate-limit) for any flow.

Elephant flow upgrade from Intelligent Application Bypass

For devices running 7.2.0 or later, you must configure elephant flow settings under the **Elephant Flow Settings** section in the access control policy (Advanced settings tab) because Intelligent Application Bypass (IAB) is deprecated from version 7.2.0 onwards for Snort 3 devices.

Post-upgrade to 7.2.0 (or later), if you are using a Snort 3 device, the elephant flow configuration settings will be picked and deployed from the **Elephant Flow Settings** section and not from the **Intelligent Application Bypass Settings** section, so if you have not migrated to Elephant Flow configuration settings, your device will lose the elephant flow configuration upon the next deployment.

This table shows the IAB or elephant flow configurations that can be applied to version 7.2.0 or later and to version 7.1.0 or earlier that are running Snort 3 or Snort 2 engines.

Cloud-Delivered Firewall Management Center	Firewall Threat Defense	Elephant Flow or IAB Configuration
Cloud-Delivered Firewall Management Center 7.0 or 7.1	Snort 2 device	Configuration from IAB is applicable.
	Snort 3 device	Configuration from IAB is applicable.
Cloud-Delivered Firewall Management Center 7.2.0	Snort 2 device	Configuration from IAB is applicable.
	Snort 3 device (7.1.0 and earlier)	Configuration from IAB is applicable.
	Snort 3 device (7.2.0 and later)	Configuration from Elephant Flow is applicable.

Configure elephant flow detection

Configure elephant flow detection to take actions on elephant flows. This helps resolve issues such as system duress, high CPU utilization, packet drops and related performance problems.

Attention Elephant flow detection is not applicable for prefiltered, trusted, or fast-forwarded flows, which do not process through Snort. As elephant flows are detected by Snort, elephant flow detection is not applicable for encrypted traffic.

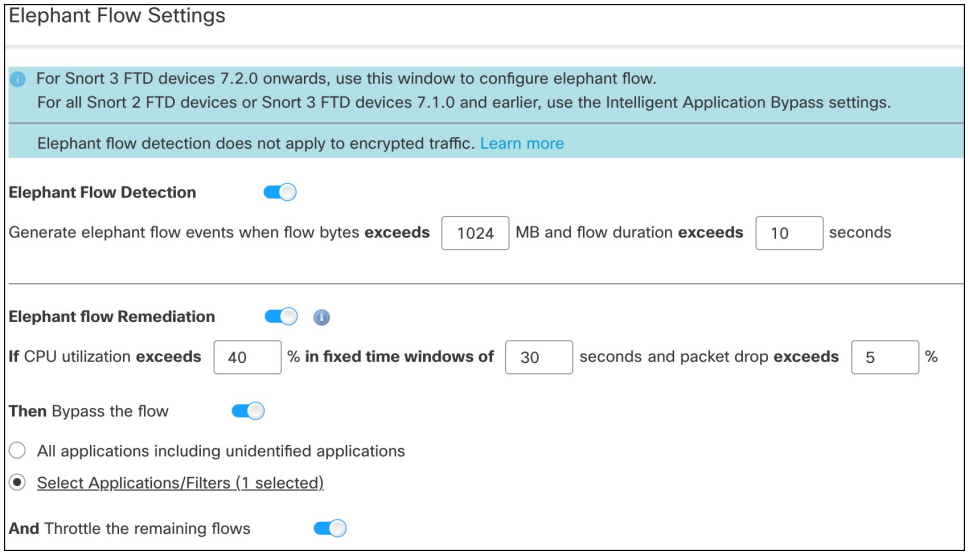
Procedure

Step 1 In the access control policy editor, click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line. Then, click **Edit** ( ) next to **Elephant Flow Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Figure 393: Configure elephant flow detection

Figure 394: Configure elephant flow detection



Step 2 Configure the values for flow bytes and flow duration for the **Elephant Flow Detection** toggle button that is enabled by default.

When they exceed your configured values, elephant flow events are generated.

Step 3 To remediate elephant flows, enable the **Elephant Flow Remediation** toggle button.

Step 4 To set the criteria for remediation of the elephant flow, configure the values for CPU utilization percentage, duration of fixed time windows, and packet drop %.

CPU utilization is calculated per elephant flow and is derived from the flow latency. If the CPU utilization crosses the configured threshold and other configurations, such as fixed time windows and packet drops, are also matched, the elephant flow remediation actions are applied. Similarly, packet drop calculation is based on the packets dropped per CPU. After the packet drop percentage exceeds the configured value on a specific CPU, the remediation actions are applied. For example, consider that configurations are set to the default, that is, CPU utilization of 40%, fixed time window of 30 seconds, and packet drop of 5%. On a specific CPU, if more than 5% of packet drops are detected and the CPU utilization per flow exceeds 40% in the fixed time frame of 30 seconds, then the flows are either bypassed or throttled.

Step 5 Configure the actions for elephant flow remediation when it meets the configured criteria.

a. **Bypass the flow**—Enable this button to bypass Snort inspection for selected applications or filters. Choose from:

   - **All applications including unidentified applications**—Select this option to bypass all the application traffic. If you configure this option, you cannot configure the throttle action (rate-limit) for any flow.
   - **Select Applications/Filters**—Select this option to select the applications or filters whose traffic you want to bypass; see the topic *Configuring Application Conditions and Filters* in the *Access Control Rules* chapter in the Cisco Secure Firewall Management Center Device Configuration Guide.

b. **Throttle the flow**—Enable this button to apply rate-limit to the flow and continue to inspect flows. Note that you can select the applications or filters to bypass Snort inspection and throttle the remaining flows.

Note
Automatic removal of throttle from a throttled elephant flow occurs when the system is out of duress, that is, the percentage of Snort packet drops is less than your configured threshold. Consequently, rate limiting is also removed.

You can also manually remove throttling from a throttled elephant flow, using these threat defense commands:

- **clear efd-throttle <5-tuple/all> bypass**—This command removes throttling from the throttled elephant flow and bypasses Snort inspection.
- **clear efd-throttle <5-tuple/all>**—This command removes throttling from the throttled elephant flow and Snort inspection continues. Elephant flow remediation is skipped after using this command.

  For more information about these commands, refer to the Cisco Secure Firewall Threat Defense Command Reference.

Step 6 In the **Remediation Exemption Rule** section, click **Add Rule** to configure L4 access control list (ACL) rules for flows that must be exempted from remediation.

a) In the **Add Rule** window, use the **Networks** tab to add the network details, that is the source network and the destination network. Use the **Ports** tab to add the source port and the destination port.

   If an elephant flow is detected and it matches the rules that are defined, an event is generated with the reason as **Elephant Flow Exempted** in the **Reason** column header of **Connection Events**.

b) In the **Remediation Exemption Rule** section, you can view the flows that are exempt from the remediation action.

c) Click **OK** to save the elephant flow settings.

Step 7 Click **Save** to save the policy.

What to do next

Deploy configuration changes.

After configuring your elephant flow settings, monitor your connection events to see if any flows are detected, bypassed, or throttled. You can view this in the **Reason** field of your connection event. The three reasons for elephant flow connections are:

- Elephant Flow
- Elephant Flow Throttled
- Elephant Flow Trusted

Attention Enabling elephant flow detection alone does not cause generation of connection events for elephant flows. If a connection event is already logged for another reason and the flow is also an elephant flow, then the **Reason** field contains this information. However, to ensure that you are logging all elephant flows, you must enable connection logging in the applicable access control rules.

Refer to Cisco Secure Firewall Elephant Flow Detection for more information.

Use case - Configure Elephant Flow Detection Outcomes

This use case provides a sample scenario for configuring outcomes and system responses that detect elephant flows in network traffic to optimize performance and resource usage.

Elephant flows

An elephant flow is an extremely large network connection that

- consists of long-running connections set up by TCP or other protocols measured over a network link,
- measures larger than 1 GB per 10 seconds by default, and
- can cause performance duress or issues in Snort cores by consuming excessive CPU resources.

Impact on system performance

Elephant flows can potentially consume an excessive amount of CPU resources and impact other competing flows for detection resources and cause issues, such as increased latency or packet drops.

Benefits of elephant flow detection and remediation

- Elephant flow configuration allows customization and the option to bypass or even throttle elephant flows.
- You can choose to bypass or throttle flows that are based on your chosen applications to provide Snort inspection of suspect traffic, while bypassing more trusted traffic.
- Elephant flow remediation helps prioritize and free up more bandwidth for your internal applications, depending on your specific requirements.

How elephant flows work

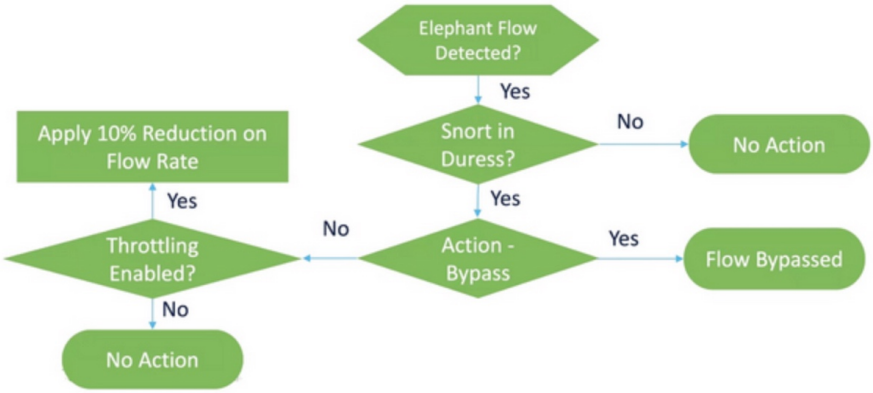
Summary

The key components involved in the elephant flow process are:

- System detection: Monitors for Snort duress conditions based on configured parameters.
- Flow identification: Determines when a flow qualifies as an elephant flow.
- CPU monitoring: Tracks CPU utilization against configured thresholds.
- Bypass action: Allows traffic for trusted applications to pass without Snort inspection.
- Throttle action: Reduces the flow throughput in 10 percent increments until CPU utilization drops below the threshold.

Workflow

Figure 395: Elephant flow workflow



These stages describe how elephant flows are processed:

1. The system continuously monitors traffic flows and applies intrusion policy processing to all flows, including potential elephant flows, ensuring that attacks are detected early in connections.
2. When an elephant flow is detected based on configured parameters, the system evaluates whether a Snort duress condition exists by checking CPU utilization and time window parameters.

   No action is taken unless the system detects a Snort duress condition (performance issue). The system does not throttle or bypass a flow just because it is large.
3. If duress conditions are met, the system determines whether to bypass or throttle the elephant flow based on configuration settings.

   The actions of throttle and bypass are mutually exclusive. This means that you can either bypass or throttle a flow, but not both.
4. For bypass actions, if configured for specific applications only, the system allows trusted application flows to pass without Snort inspection while throttling remaining flows causing duress.

   This ensures that other non-trusted application flows still receive full Snort inspection, although their bandwidth is reduced.
5. For throttle actions, the system reduces flow throughput in 10 percent increments until CPU utilization reduces to below the configured threshold.

Sample business scenarios

In a data center, several activities are happening, such as replication of data between clusters, virtual machine integration, and database backup. Users in an organization could be watching videos on an OTT service or downloading them. Bandwidth utilization for such activities might result in elephant flows that slow down the network and impact the performance of important tasks. As a network administrator (and depending on your specific requirements), you want visibility into such large flows that are causing bandwidth issues and remediate them.

As an example, let's see how you can configure elephant flow parameters to bypass Snort inspection for WebEx traffic (which your organization uses for real-time video conferencing) and throttle the remaining applications or connections, including videos, movies, and so on.

Prerequisites

Ensure that you are running management center 7.2.0 or later and that the managed threat defense is also 7.2.0 or later.

When you enable only elephant flow detection, no additional connection events are generated. The system adds the Elephant Flow notation to matching connections that are already logged to the management center. **To log these events, you must enable connection logging in your access control policy.** You can do that for specific rules or add a Monitor rule that logs all connections, including elephant flows.

Configure elephant flow parameters

Configure elephant flow detection and remediation settings to identify high-bandwidth connections, and manage them through bypassing or throttling to maintain optimal system performance during high CPU utilization.

Elephant flows are high-bandwidth, long-duration connections that can impact system performance. By configuring elephant flow parameters, you can detect these flows and apply appropriate remediation actions such as bypassing trusted applications from inspection or throttling flows that cause system duress.

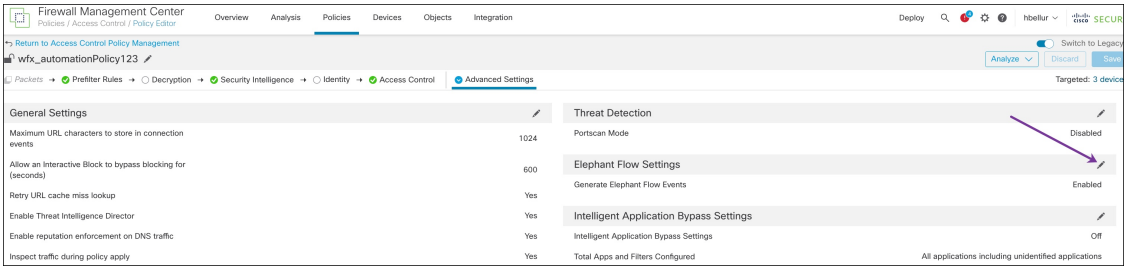
Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy that you want to edit.

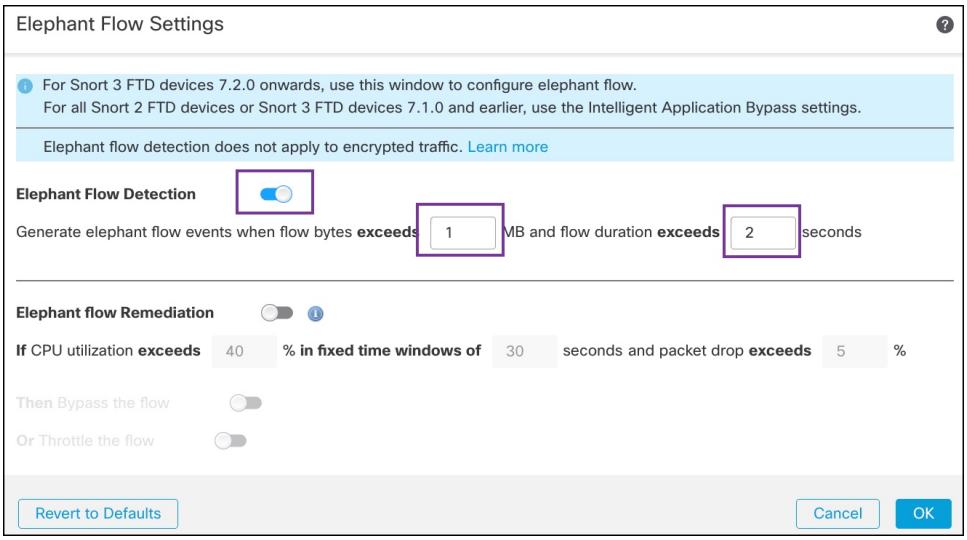
Step 3 Choose **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 Click **Edit** (✎) next to **Elephant Flow Settings**.



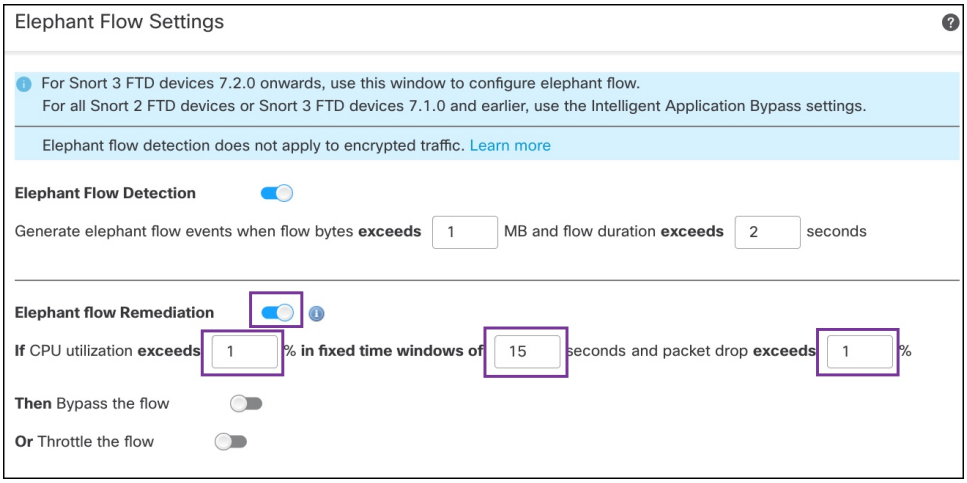
Step 5 The **Elephant Flow Detection** toggle button is enabled by default. The default setting enables detection only and no default action is configured. The detection settings allow you to adjust the flow bytes and duration so that you can identify the elephant flows in your system.

As a test setting, configure the flow bytes and duration parameters, as shown in the following figure.

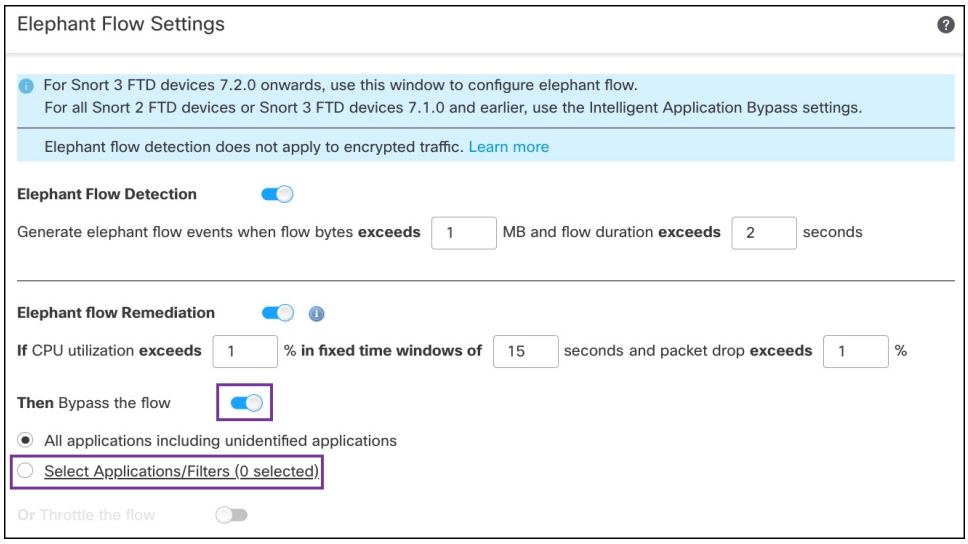


Step 6 Enable the **Elephant Flow Remediation** toggle button. When an elephant flow is detected, you can choose to bypass or throttle the flow. Bypassing a flow means that the traffic is allowed to pass without Snort inspection. Throttling indicates that the flow throughput is reduced. This rate reduction is done in 10 percent increments until the CPU utilization decreases to less than the configured threshold.

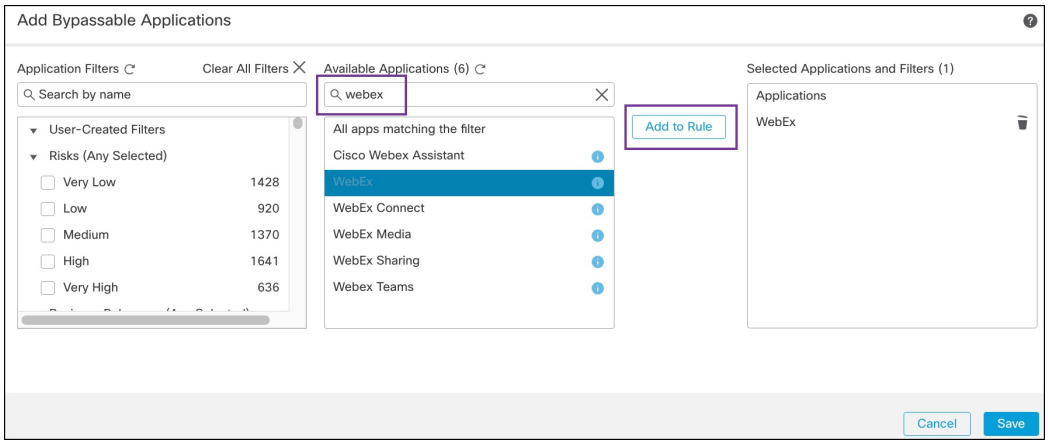
As a test setting, configure the elephant flow remediation parameters as shown in the following figure.



Step 7 Enable the **Bypass the flow** toggle button and click the **Select Applications/Filters** radio button.



Step 8 Under **Application Filters**, search for and select the **WebEx** application, add it to the rule, and click **Save**. This means that WebEx connections are trusted and prioritized. These connections will skip Snort inspection if detected as elephant flows based on the configured parameters.



Step 9 Enable the **Throttle** toggle button to throttle the remaining flows (causing duress). This ensures that all the other flows are slowed down in 10 percent increments until the Snort duress condition is met.

Click **OK** and then **Save**.

What to do next

Deploy configuration changes.

View events for elephant flows

This task enables you to monitor elephant flow connection events to track flow detection, bypass, and throttling activities through the Reason field in connection events.

After configuring your elephant flow settings, monitor your connection events to see if any flows are detected, bypassed, or throttled. You can see this information in the **Reason** field of your connection events. The three types for elephant flow connections are:

- Elephant Flow
- Elephant Flow Throttled
- Elephant Flow Trusted

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 On the **Events** tab, in the search field, type "Reason" and then "Elephant Flow" to search for elephant flow events.

Tip
To see **Elephant Flow Trusted** or **Elephant Flow Throttled** events, type "Reason" and then "Elephant Flow Trusted" or "Elephant Flow Throttled", depending on what you want to search.

Step 3 View the elephant flow that was detected mid-flow and the **Reason** field shows **Elephant Flow**. At the end of the flow, it was bypassed and the **Reason** field shows **Elephant Flow Trusted**.

Configure elephant flow remediation exemption

This task allows you to create exemption rules for specific flows that should not be subject to remediation actions even when detected as elephant flows.

You can configure L4 access control list (ACL) rules for flows that must be exempted from remediation. If a flow is detected as an elephant flow and it matches the rules that are defined, that flow is exempted from remediation.

Before you begin

You must be running management center version 7.4.0 or later and the managed threat defense must also be on version 7.4.0 or later.

Follow these steps to configure elephant flow remediation exemption.

Procedure

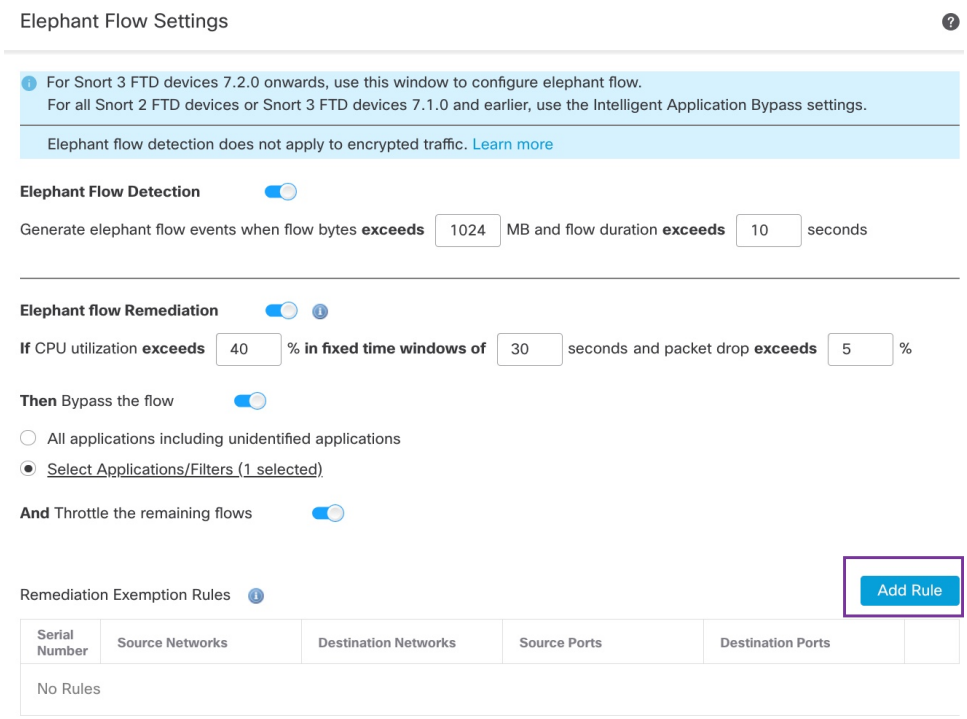
Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** ( ) next to the access control policy you want to edit.

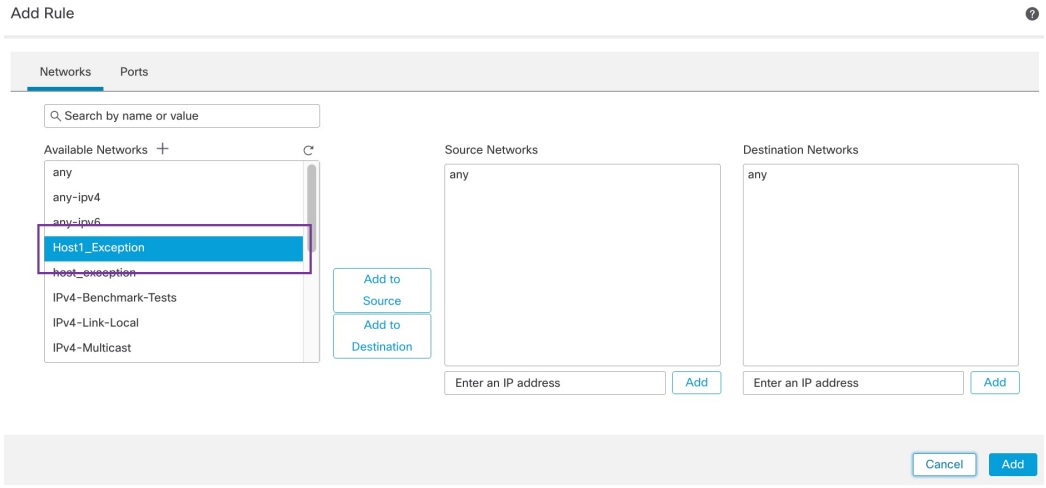
Step 3 Choose **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 Click **Edit** ( ) next to **Elephant Flow Settings**. Ensure that you have configured the elephant flow detection and remediation parameters. See Configure elephant flow parameters, on page 2741.

Step 5 Click the **Add Rule** button next to **Remediation Exemption Rules**.



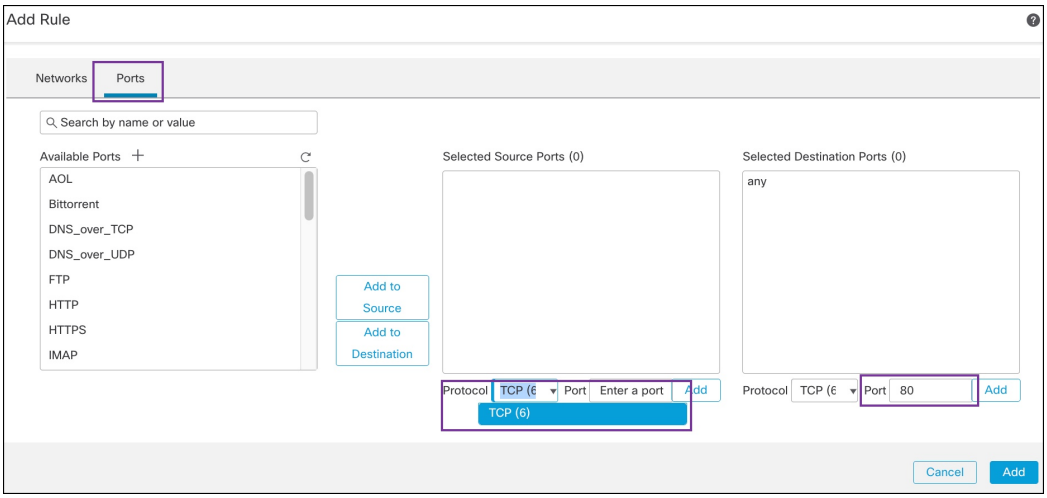
Step 6 From the list of **Available Networks**, choose the configured host to exempt from elephant flow remediation. For the purposes of this example, we have created a host called "Host1_Exception."



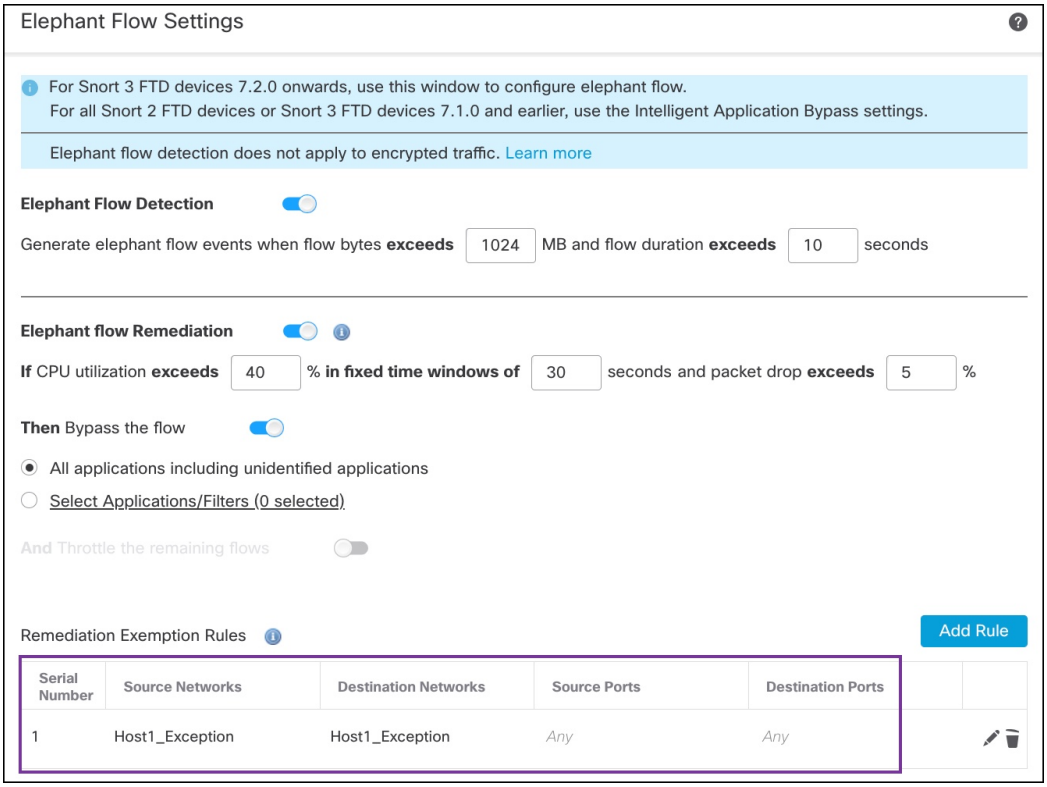
Step 7 Click **Add to Source** or **Add to Destination** as needed to add this host to the source or destination.

Step 8 Click the **Ports** tab.

a) For the source port, choose **Protocol** as TCP, enter **80** as the destination port, then click **Add**.



b) Click **OK**.



Step 9 Click **Save**.

What to do next

Deploy configuration changes.

View events for elephant flow remediation exemption

This task allows you to monitor and review elephant flows that have been exempted from remediation to understand network traffic patterns and exemption reasons.

Elephant flows are high-volume network connections that may require special handling. When these flows are exempted from remediation, you can view the exemption events to track and analyze the decision.

Procedure

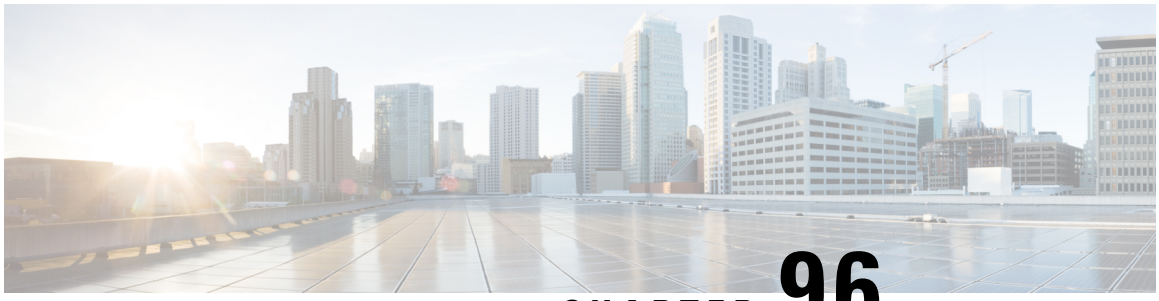
Step 1 Choose **Analysis** > **Unified Events**.

Step 2 View the elephant flows that were exempted from remediation. The **Reason** field shows **Elephant Flow Exempted**.

Additional references

For detailed conceptual information, see the Elephant Flow Detection for Snort 3 chapter in this guide or the content in the following link:

- Elephant Flow Detection



CHAPTER 96

Policy Analyzer and Optimizer

Security Cloud Control provides the Policy Analyzer and Optimizer, which is an intelligent cloud service that can analyze security policies, detect anomalies, and provide recommendations on remediations that can be performed to optimize the policies.

- About Policy Analyzer and Optimizer, on page 2749
- Common Policy Analyzer and Optimizer concepts, on page 2750
- Prerequisites to use Policy Analyzer and Optimizer, on page 2751
- Enable Policy Analyzer and Optimizer for Security Cloud Control-managed On-Premises Firewall Management Center, on page 2752
- Policy Analyzer and Optimizer licensing requirements, on page 2753
- Open Policy Analyzer and Optimizer and select a data source, on page 2753
- Search and filter policies in Policy Analyzer and Optimizer, on page 2754
- Access control policy analysis and optimization , on page 2754
- Network Address Translation policy analysis, on page 2764
- Troubleshooting Policy Analyzer and Optimizer, on page 2770
- Frequently Asked Questions About Policy Analyzer and Optimizer, on page 2771

About Policy Analyzer and Optimizer

AgenticOps for firewalls provides operational insights and tools that help improve firewall management, policy hygiene, and deployment readiness. One such capability is Policy Analyzer and Optimizer, which uses an algorithmic approach to evaluate firewall policy rules and identify anomalies such as duplicate, shadowed, redundant, or otherwise optimizable rules. For more information, refer to AgenticOpsInsights to know more about the various other functionalities that AgenticOps provides.

Secure Firewall Threat Defense devices with extensive policies may have numerous duplicate or shadowed rules. Such large policies with unoptimized rulesets can lead to excessive consumption of device memory, delayed loading of rules, and long search duration, resulting in inefficient security policy enforcement, reduced network speeds, and extended deployment durations.

Policy Analyzer and Optimizer is an intelligent Security Cloud Control Firewall Management service that analyzes firewall policies, detects rule anomalies, and helps you understand where a policy can be optimized. Policy Analyzer and Optimizer supports access control policies policies for Cloud-Delivered Firewall Management Center and supported Security Cloud Control-managed On-Premises Firewall Management Center data sources.

In addition, Policy Analyzer and Optimizer can do the following:

- View policy health and optimization opportunities for the selected management center data source.
- Analyze policies on demand or rely on scheduled analysis that runs every 24 hours.
- Download analysis reports as PDFs after analysis completes.
- Use Access Control Policy Analyzer and Optimizer remediation workflows where supported.
- Stage and apply remediation for an entire anomaly category, selected observations, or individual rules within supported observations.
- Use Network Address Translation Policy Analyzer and Optimizer findings and AgenticOps insights to investigate NAT policies that contain optimizable rules.

Supported policy types

Security Cloud Control Firewall Management performs anomaly analysis on the following policy types:

Table 147: Supported policy types for Policy Analyzer and Optimizer

Policy type	Use this section when	Current scope
Access control policy	You want to analyze access control policies, review access control findings, and apply supported remediations.	Analysis, remediation, reporting, remediation history, and policy insights.
Network address translation policy	You want to analyze NAT policies for shadowed and redundant NAT rules.	Analysis and reporting for fully shadowed and fully redundant NAT rules. NAT remediation is not available in this release.

Common Policy Analyzer and Optimizer concepts

Policy Analyzer and Optimizer uses these shared concepts across access control policy analysis and NAT policy analysis.

Table 148: Common Policy Analyzer and Optimizer concepts

Concept	Common behavior	Policy-type difference
Data source	A Cloud-Delivered Firewall Management Center or an On-Premises Firewall Management Center selected in Policy Analyzer and Optimizer.	The Access Control tab can be used for supported On-Premises Firewall Management Center data sources. The Network Address Translation tab is available only when Network Address Translation Policy Analyzer and Optimizer is enabled through AgenticOps for the selected data source.

Concept	Common behavior	Policy-type difference
Sync	Fetches current policy inventory and metadata, including newly created, deleted, or modified policies.	Sync does not, by itself, create a new analysis summary. If the policy changed after the last analysis, run analysis again or wait for scheduled analysis.
Analysis	Evaluates selected policies and produces policy health information, observations, and report data.	Access control Policy Analyzer and Optimizer evaluates access control rule and object findings. Network Address Translation Policy Analyzer and Optimizer evaluates NAT rule shadowing and redundancy.
Reporting	Provides a downloadable analysis report after analysis completes.	Access control Policy Analyzer and Optimizer also provides remediation reports after supported remediation is applied.
Remediation	A workflow that stages selected changes and applies them only when you select Apply Remediation.	Available for access control policies. Network Address Translation Policy Analyzer and Optimizer does not provide an Apply Remediation workflow. You can stage remediation for an entire anomaly category, selected observations within a category, or individual rules within an expanded observation. You can also select individual rules within an expanded observation. Using the suggested remediations, you can disable or delete selected Duplicate Rules and Expired Rules, remove selected fully overlapped objects from rules, and merge selected Mergeable Rules.

Prerequisites to use Policy Analyzer and Optimizer

Before you use Policy Analyzer and Optimizer, verify the prerequisites that apply to your policy type and data source.

- The On-Premises Firewall Management Center must be integrated with Security Cloud Control.

Note In the On-Premises Firewall Management Center, navigate to **Integration** > **Cisco Security Cloud** and check the **Enable Policy Analysis & Optimization** checkbox after integrating with the Cisco Security Cloud.

- The Cloud-Delivered Firewall Management Center must be provisioned in Security Cloud Control Firewall Management.

Prerequisites to use Policy Analyzer and Optimizer for access control policy

- Access control policy analysis is supported on the On-Premises Firewall Management Center operating on Version 7.2 or later. The Cisco-recommended Version is 7.6 or later, because these versions support cross-launch to Policy Analyzer and Optimizer directly from On-Premises Firewall Management Center.
- Access control policy to be analyzed must be associated with a Firewall Threat Defense device that is managed by an On-Premises Firewall Management Center integrated with Security Cloud Control.

Prerequisites to use Policy Analyzer and Optimizer for Network Address Translation

- NAT policy analysis is supported on the On-Premises Firewall Management Center operating on Version 7.6 or later.
- Network Address Translation Policy Analyzer and Optimizer availability also depends on AgenticOps being enabled for the selected tenant and data source. For more information, refer to Enable AgenticOps Insights.
- Enable AgenticOps for the Cloud-Delivered Firewall Management Center or On-Premises Firewall Management Center data source that you want to analyze. Choose **Insights & Reports** > **Settings** > **Policy Analyzer and Optimizer** > **Data Sources** and enable AgenticOps for the management center you want.

Note If an On-Premises Firewall Management Center is running a version earlier than 7.6, access control Policy Analyzer and Optimizer may still be available for analyzing access control policy, but Network Address Translation Policy Analyzer and Optimizer is not available for that data source.

Enable Policy Analyzer and Optimizer for Security Cloud Control-managed On-Premises Firewall Management Center

Onboard your On-Premises Firewall Management Center to Security Cloud Control, navigate to **Administration** > **Integrations** > **Firewall Management Center**, select the On-Premises Firewall Management Center, and choose **Policy Analyzer and Optimizer** under **System** in the right pane. See Onboard an On-Premises Firewall Management Center for more information.

If you have an On-Premises Firewall Management Center Version 7.6 and want to use Policy Analyzer and Optimizer, follow the steps below:

Procedure

Step 1 In your On-Premises Firewall Management Center, navigate **Integration** > **Security Cloud Control**.

Step 2 If you have not integrated your On-Premises Firewall Management Center with Cisco Security Cloud, click **Enable Security Cloud Control** and follow the steps. To authorize the cloud integration, you must choose an existing Security Cloud Control tenant or provision a new one, to which your on-premises Firewall Management Center will get onboarded, after the cloud integration is successful.

Step 3 After integrating your On-Premises Firewall Management Center with Security Cloud Control, check the **Enable Policy Analyzer and Optimizer** checkbox and click **Save**.

Policy Analyzer and Optimizer licensing requirements

Policy Analyzer and Optimizer requires no additional licensing. It comes as part of the Security Cloud Control Firewall Management base subscription.

Open Policy Analyzer and select a data source

You can open Policy Analyzer and Optimizer from Security Cloud Control Firewall Management or cross-launch from supported On-Premises Firewall Management Center policy pages.

Follow these steps to open Policy Analyzer and Optimizer and select a data source:

Procedure

Step 1 In the Security Cloud Control Firewall Management left pane, choose **Administration** > **Integrations** > **Firewall Management Center**.

Step 2 Select **Cloud-Delivered FMC** or an On-Premises Firewall Management Center.

Step 3 Click Policy Analyzer and Optimizer under **System**.

Note
Alternatively, choose **Insights & Reports** > **Policy Analyzer and Optimizer**.

Step 4 Use the data source selector at the top of the Policy Analyzer and Optimizer page to choose the Cloud-Delivered Firewall Management Center or On-Premises Firewall Management Center whose policies you want to review.

Step 5 Choose the **Access Control** tab depending on the policy type that you want to analyze.

Note
If the selected data source is not enabled for AgenticOps, the **Network address translation** tab shows a message indicating that the NAT Policy Analyzer and Optimizer require AgenticOps for the selected data

source. However, you can still utilize the access control Policy Analyzer and Optimizer if the access control prerequisites are met.

Sync and analysis behavior

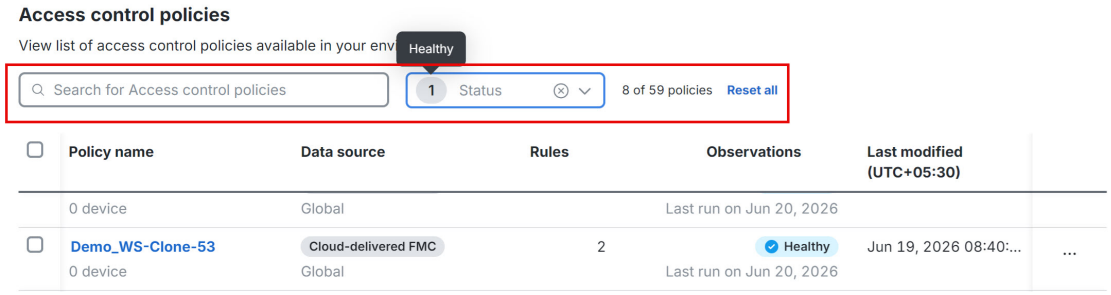
Policy Analyzer and Optimizer displays policy metadata, including analysis status, last modified time, and last analyzed time. Last modified reflects changes made to the policy in the management center that manages it. Last analyzed reflects the analysis run that produced the current Policy Analyzer and Optimizer summary.

Use **Sync with FMC** to fetch the latest access control and NAT policy inventories and modified timestamps from the management center. Each Policy Analyzer and Optimizer tab shows the sync success or failure status for the policy type on that tab. If a policy was modified after the last analysis, the displayed analysis can be out of date until you run analysis again or a scheduled analysis runs.

Search and filter policies in Policy Analyzer and Optimizer

Use the **Search for access control policies** field and the **Status** field to narrow the access control policy list.

Search helps you find a policy by name. Status filtering helps you find policies by analysis state, such a **Healthy**, **Error**, **In progress**, **Analysis pending**, **Optimizable**, or **Analysis out of date**.



- **Search for access control policies**: Finds access control policies that match the policy name you enter.
- **Status**: Filters access control policies by their current analysis state.

Access control policy analysis and optimization

Use the **Access Control** tab when you want to analyze access control policies, review access control policy findings, and apply the remediations that Policy Analyzer and Optimizer supports for access control rules and objects.

After provisioning a Cloud-Delivered Firewall Management Center or onboarding an On-Premises Firewall Management Center to your Security Cloud Control Firewall Management tenant and creating policies, you can start analyzing them using the Policy Analyzer and Optimizer.

For more information, refer to Onboard an On-Premises Firewall Management Center and Enable Cloud-Delivered Firewall Management Center on Your Security Cloud Control.

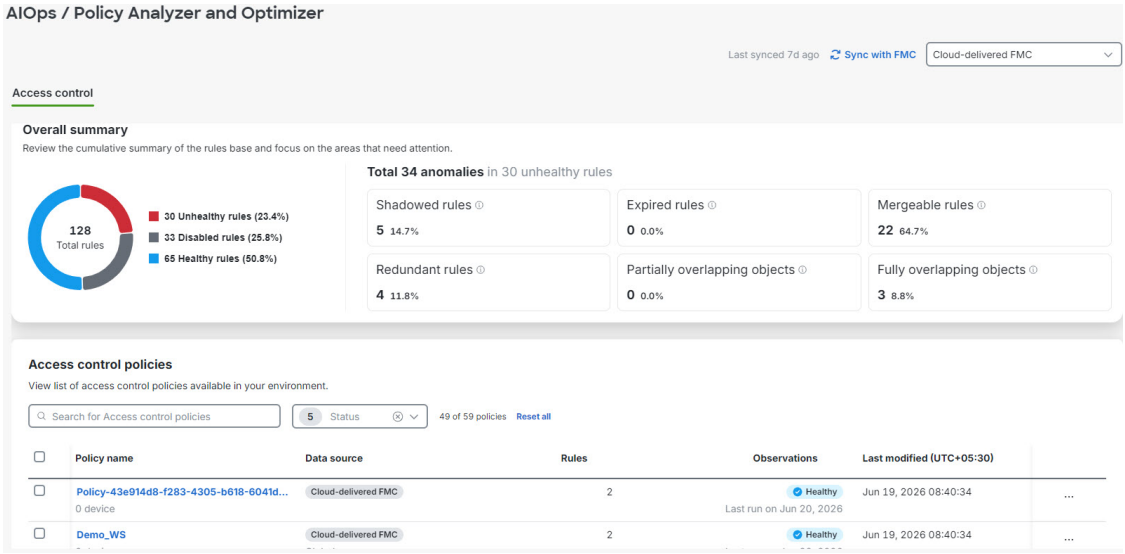
Analyze access control policies

Use this procedure to assess an access control policy before making changes, review policy health, or re-run analysis after policy updates. If a policy has not been analyzed, you can start a new analysis. If the analysis is out of date, you can re-analyze the policy to refresh the results.

Note When you create a new policy, it might take a while for the Policy Analyzer and Optimizer to fetch the policy details and show up on the **Policy Analyzer and Optimizer**. Click the refresh ( ) button on the top-right corner to manually refresh the page to see new policies.

Procedure

Step 1 Choose **Insights & Reports** > **Policy Analyzer and Optimizer**.



Step 2 In the right pane, select **Cloud-delivered FMC** or an On-Premises Firewall Management Center from the drop-down list as the data source whose policies you wish to analyze.

Step 3 In the **Access Control** area, select the policy you want to analyze or review.

Note

- For an unanalyzed policy, click on **Analyze Policies**.
- If the policy status indicates **Analysis out of date**, click on **Re-analyze Policy** under **Analysis Actions** on the right.

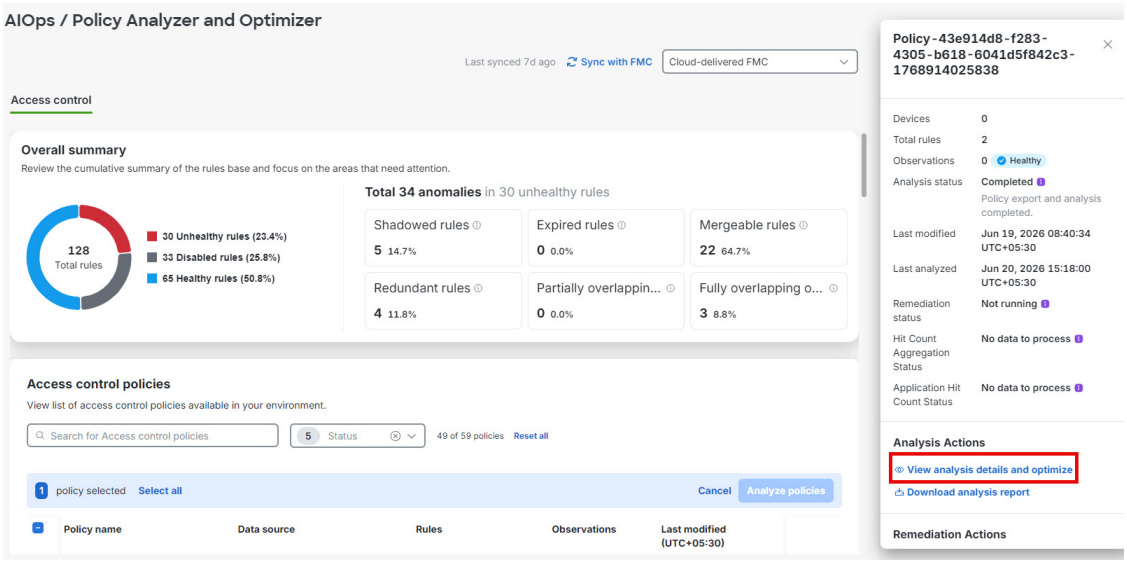
The **Overall summary** section displays the total number of rules categorized by their health status: healthy, disabled, or unhealthy, for the selected management center (Cloud-Delivered Firewall Management Center or On-Premises Firewall Management Center).

The dashboard also highlights specific anomalies within your unhealthy rules. You can review the count and percentage for the categories: Shadowed rules, Expired rules, Mergeable rules, Redundant rules, Partially overlapping rules, Fully overlapping objects.

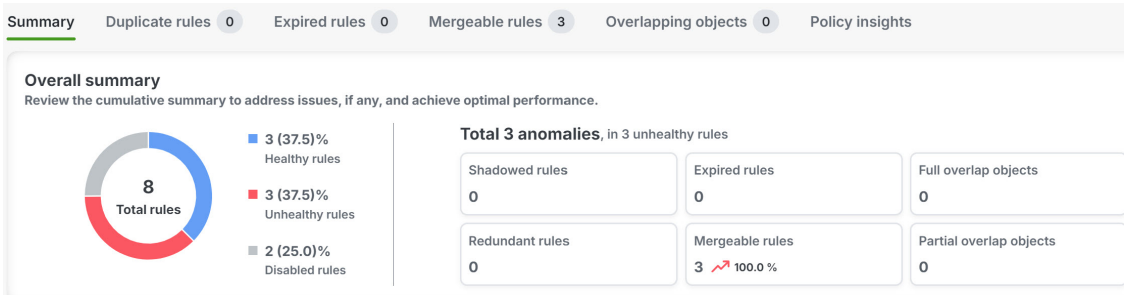
Access control policy analysis summary

When analysis completes, select an access control policy to review the right-pane summary and click **View analysis details and optimize**.

Alternatively, you can click the ellipse button associated with a policy and select **View details** to see the right-pane summary.



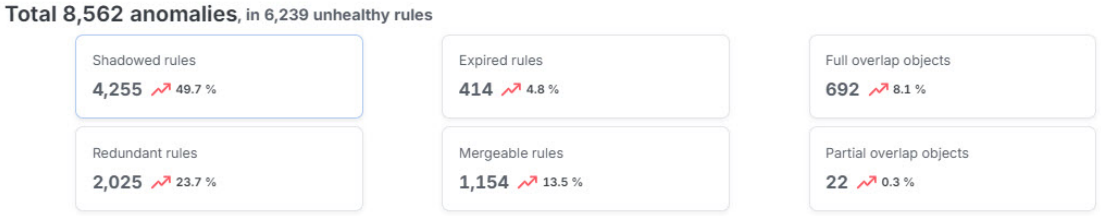
Overall summary—provides insights on how many rules are healthy, disabled, expired, and contain anomalies, using a pie chart for the selected access control policy. You can also hover over the part of the pie to view the percentage of rules.



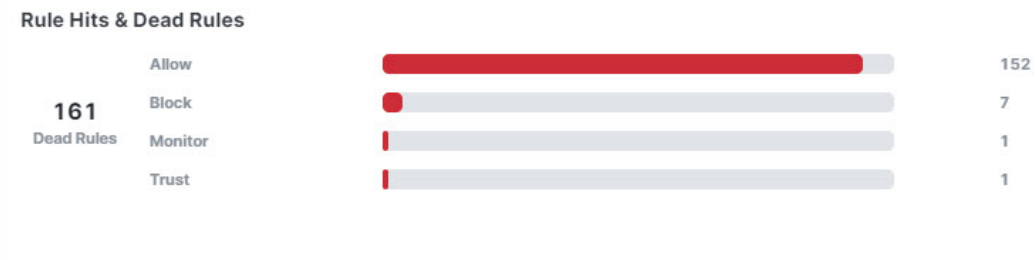
Rule usage history—shows how recently rules were used, with time periods.



Rules with Anomalies—provides insights on how many rules have anomalies.



Hits rules & dead rules—provides insights on hitcount of expired rules, for rule types including allow, block, monitor, and trust.

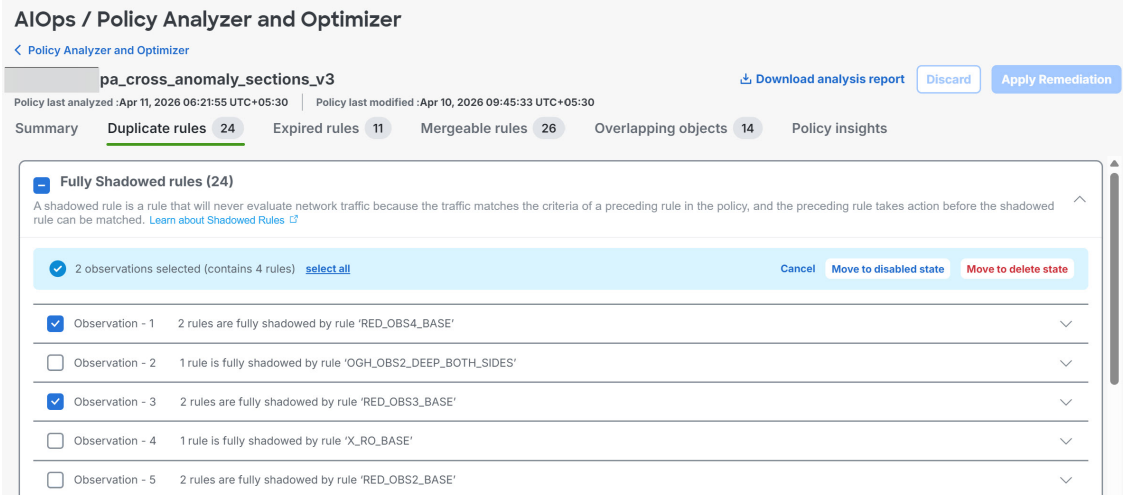


In remediation tabs such as **Duplicate rules**, **Expired rules**, **Mergeable rules**, **Overlapping objects**, and **Policy insights** you can select the check box next to a category to stage all observations in that category, or expand the category and select specific observations or rules. Available actions depend on the anomaly type.

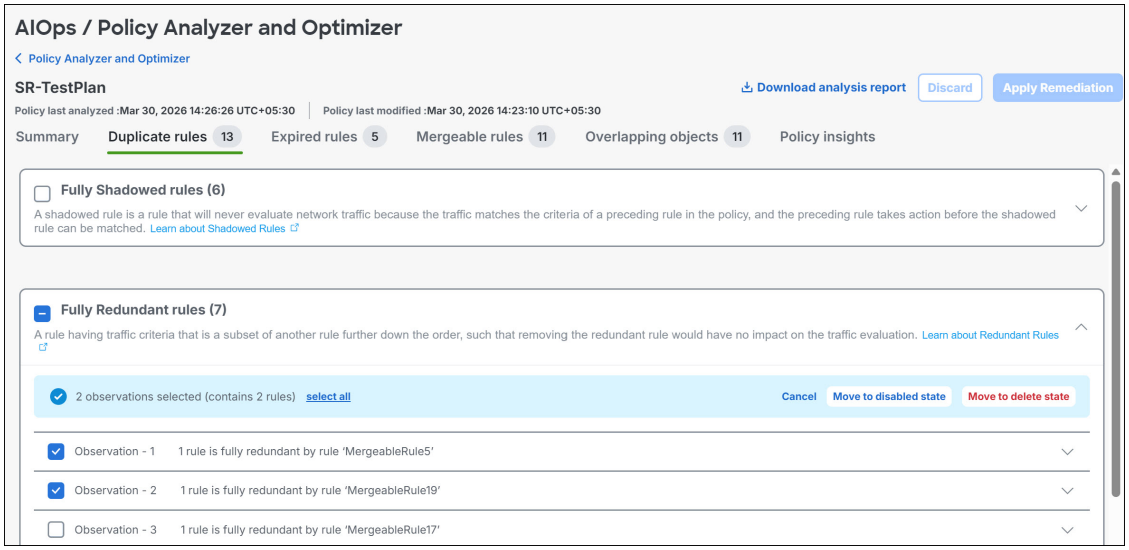
Analyze duplicate rules in access control policy

The **Duplicate Rules** tab lists shadowed and redundant rules with anomalies:

- A **Fully Shadowed rules** is one that will never evaluate network traffic because another rule that precedes it over shadows this rule.



- A **Fully Redundant rules** is one that is just a part of another larger rule, such that removing this redundant rule does not have an impact on the network traffic, because the traffic evaluation that this rule must perform is already performed by another rule.



You can remediate all duplicate-rule observations in a category, selected fully shadowed or fully redundant observations, or individual rules within an expanded observation. For selected duplicate rules, choose **Move to disabled state** or **Move to delete state**. Disable rules first when you want to measure the impact before deleting them.

Note Expand each observation to review the affected rules before you stage a remediation. Each rule in the list is displayed with a set of attributes; click the settings button on the top right to select which rule attributes you want to display along with the rule.

When duplicate-rule remediation removes shadowed or redundant rules and retains a base rule, Policy Analyzer and Optimizer adds a standardized comment to the retained rule. The comment identifies the retained rule and the removed rules, which helps you audit the cleanup.

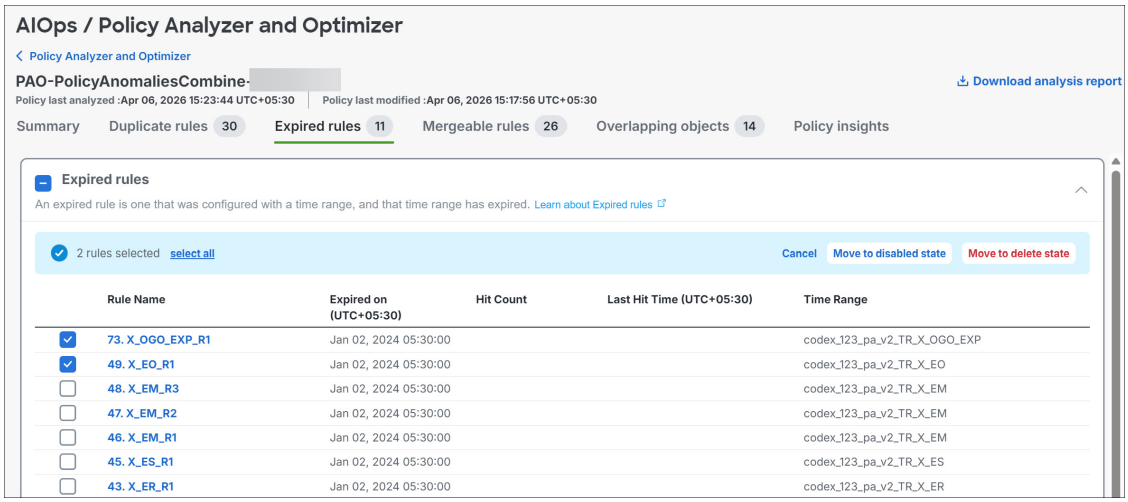
After you stage the selected duplicate-rule remediations, you can still **Undo** them before clicking **Apply Remediation**. It is recommended that you disable rules first to measure the impact and delete them later, because deleting them permanently removes them.

You can enable the disabled rules any time by navigating to the Cloud-Delivered Firewall Management Center or the On-Premises Firewall Management Center on which the rules are present.

Analyze expired rules in access control policy

The **Expired Rules** tab lists rules that were configured with a time range and the time range has expired. You can also see rule information such as the date on which the rule expired, hit count, last hit time, and the time range.

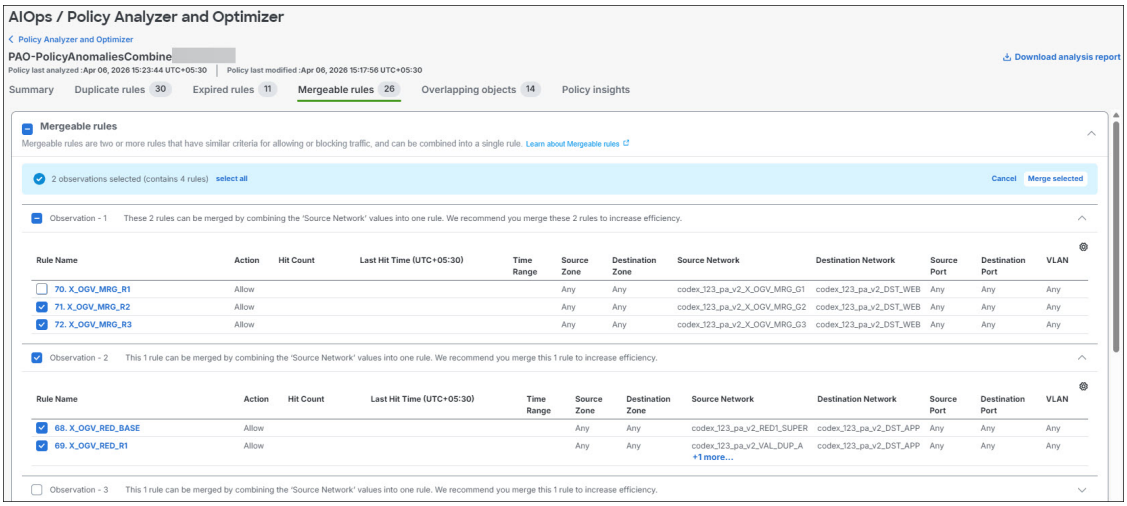
You can remediate all expired-rule observations in the category, selected expired-rule observations, or selected expired rules. For selected expired rules, choose **Move to disabled state** or **Move to delete state**. Only selected rules are changed when you apply remediation.



Analyze mergeable rules in access control policy

The **Mergeable Rules** tab lists the rules that have similar allow and block settings and can be merged into a single rule.

Review the mergeable-rule observations and stage remediation for the entire category, selected observations, or individual rules within an expanded observation. Click **Merge Selected** to merge only the staged items. Unselected mergeable-rule observations remain unchanged until you stage and apply remediation for them.

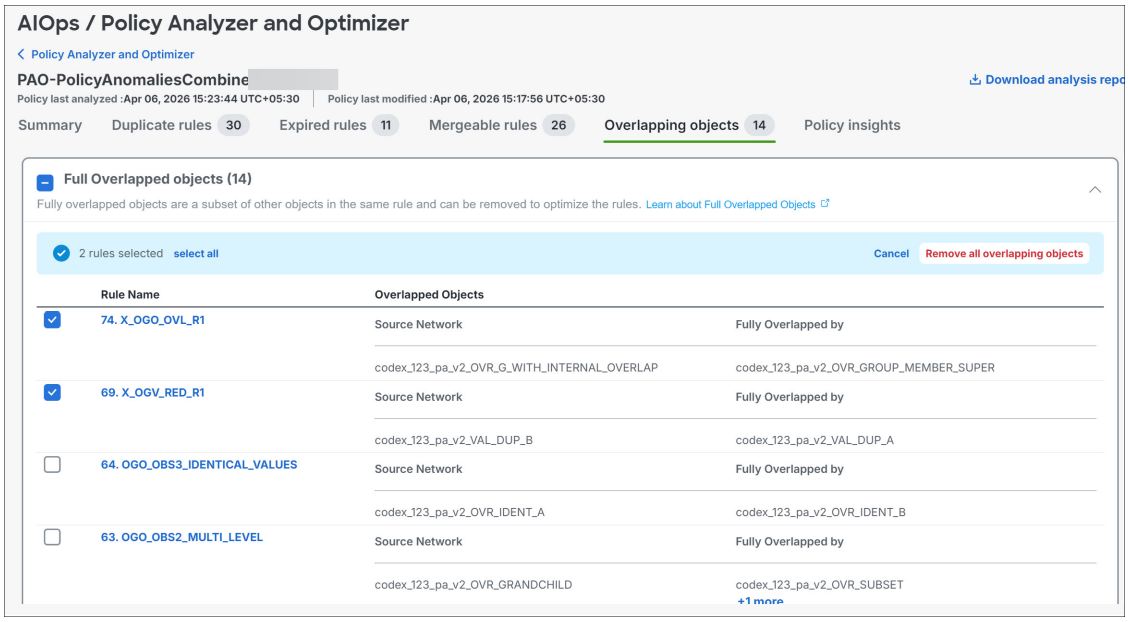


Note When you select individual rules to merge, select only consecutive rules. The first selected rule in the sequence is updated with the merged rule criteria, and the remaining selected rules are removed. You can start the merge from any rule in a consecutive set. The first rule in the observation does not need to be selected.

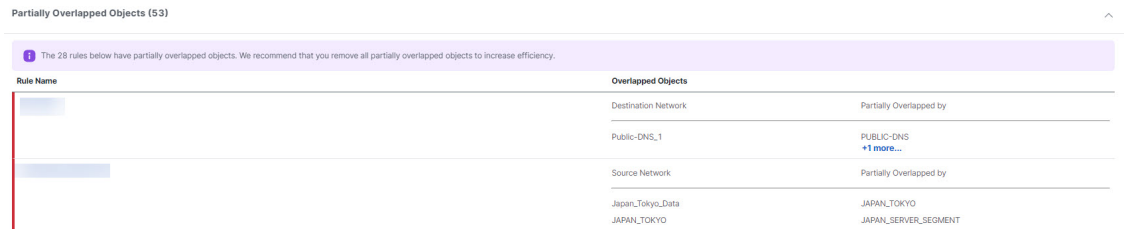
Analyze overlapping objects in access control policy

The **Overlapping Objects** tab lists objects that are either fully overlapping (the IP addresses or port numbers are either the same or a complete subset) or partially overlapping (some subset of IP addresses are repeated, but not all).

For example, if a rule contains an object for 192.168.1.1 and another for 192.168.1.0/24, the 192.168.1.1 object is fully overlapped by the other object and is not needed in the rule. For fully overlapped objects, you can remediate all fully overlapped observations in the category, selected observations, or selected rules. Click **Remove All Fully Overlapped Objects from Rules** to remove fully overlapped objects only from the staged items. For partial overlaps, evaluate each occurrence and edit the objects directly.



For partial overlaps, you need to evaluate each occurrence, determine if any changes can be made, and implement those changes directly by editing the objects.

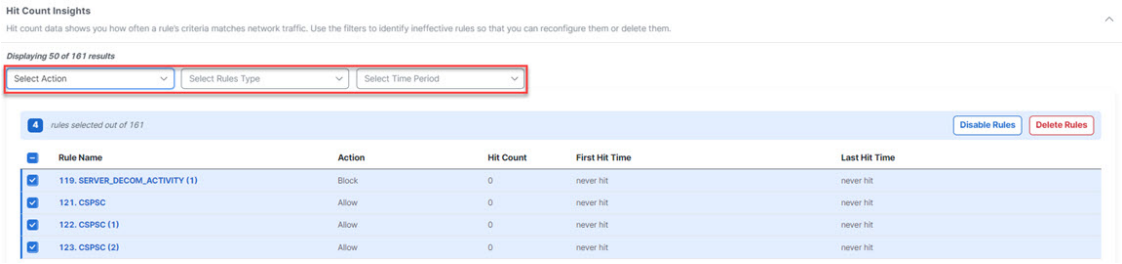


Analyze policy insights in access control policy

The **Policy Insights** tab has a **Hit Count** section that initially lists any rules that have never been triggered (**Never Hit Rules**). The hit count information is from all devices that are assigned to the policy. You can change criteria and see other hit count information, for example, **Not Hit Rules** for the past six months, or **Hit Rules** over a selected time period. You can filter the rules using the actions set in the rules, hit information, and time period:

- **Never Hit Rules**—Rules that have never been hit from the time they were created.
- **Hit Rules**—Rules that have been hit in the selected time period.
- **Not Hit Rules**—Rules that have not been hit in the selected time period.

Select the rules that you want to disable or delete, and select **Disable Rules** or **Delete Rules**. These changes are staged until you select **Apply Remediation**. First, disable the rules when you want to measure their impact before deleting them.



Access control policy remediation

Policy Analyzer and Optimizer stages remediation changes before it updates the policy. You can stage remediation for an entire anomaly category, selected observations, or individual rules within an expanded observation. Before you apply remediation, unselected observations and rules remain unchanged and available for later selection in the same analysis report.

Review staged changes before you select **Apply Remediation**. After you apply remediation, you cannot apply more changes from the same analysis report. To remediate remaining anomalies, run policy analysis again on the updated policy and use the new report.

Before you begin

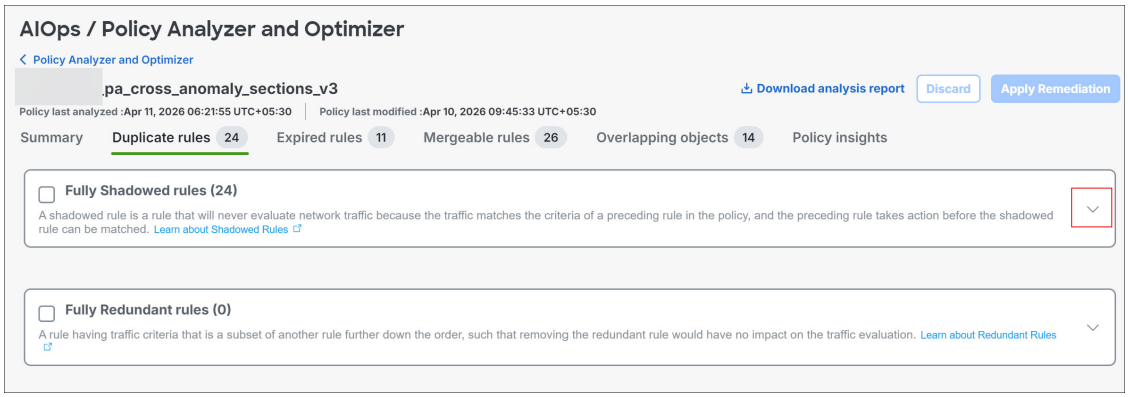
- Back up all policies before you apply remediation.
- Ensure that at least one remediation is staged. If no remediation is staged, **Apply Remediation** is disabled.
- Verify the **Policy Last Modified** and **Policy Last Analyzed** timestamps, and review the number of rules that are staged for remediation.

Procedure

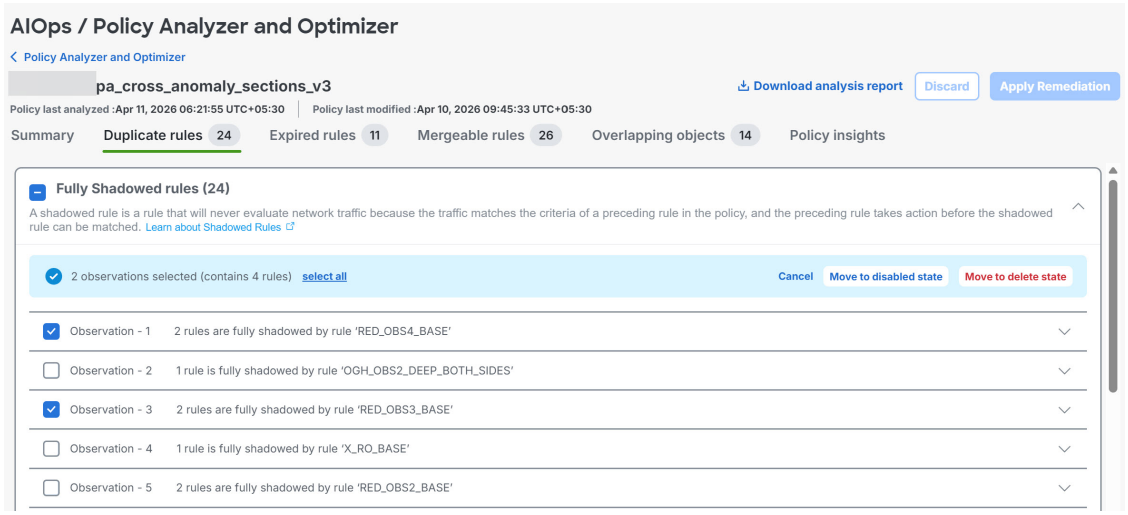
Step 1 In the **Policy Analyzer and Optimizer** page, select the policy to see details about the analysis on the right pane and click **View analysis details & optimize**.

Step 2 Click the remediation tab that contains the anomalies that you want to fix.

Step 3 Expand a remediation category. Select the check box for the category to stage all observations, or select the check boxes for specific observations or rules to stage only those items. Repeat this step in other remediation tabs, as needed.



Step 4 Select the remediation action that applies to the staged items, such as **Move to disabled state**, **Move to delete state**, **Merge Selected**, **Remove All Fully Overlapped Objects from Rules**, **Disable Rules**, or **Delete Rules**.



Step 5 Review the staged remediation. If you need to change the staged remediation before applying it, use the available undo or discard action.

Step 6 Click **Apply Remediation**.

Step 7 Read the confirmation message, which summarizes the remediations that will be applied. Confirm that the selected policy and staged remediation are correct.

Step 8 Click **Apply**.

Note
For an On-Premises Firewall Management Center in which the Change Management Workflow is enabled, when policy remediations are applied, an internal workflow ticket is created and the changes are staged. The changes take effect only when the ticket is submitted or approved. See Change Management in *Cisco Secure Firewall Management Center Administration Guide* for more information.

After remediation is complete, the selected rules are updated. Policy Analyzer and Optimizer automatically analyzes the updated policy and generates a refreshed summary that shows any remaining issues. Verify the intended changes in the corresponding access control policy.

Download access control analysis report

The access control analysis report is a PDF summary of completed remediations for a selected policy. It includes only the remediation categories that apply to the changes you selected and applied. Each section lists the rule name, remediation action, and related comments. For example, if no duplicate rules were remediated, the report excludes the duplicate-rule remediation section.

The report includes only the remediation sections and rule details for remediations that you selected and applied. Unselected anomalies are not included in the report.

1. In the Policy Analyzer and Optimizer page, click **Access control**.
2. Select an access control policy, and in the right pane, click **Download analysis report**.

The report can include these sections:

- Remediation Summary
- Hit Count Remediation
- Expired Rules Remediation
- Duplicate Rules Remediation
- Mergeable Rules Remediation



Note To determine whether a policy is remediated by the Policy Analyzer and Optimizer, navigate to **Policies** > **Access Policies** and edit a policy to view the rules in the **Policy Editor**. When a policy is remediated, a comment is added to the rules that are optimized.

You can also filter all the optimized rules using "updated by Policy Analyzer and Optimizer" to view all the remediated rules.

During duplicate rule remediation, the retained rule is updated with a comment specifying which rule was preserved and which were removed.

Network Address Translation policy analysis

Use the **Network Address Translation** (NAT) tab when you want to analyze NAT policies and find rules that are shadowed or redundant. NAT Policy Analyzer and Optimizer helps you identify rules that are not contributing to effective NAT behavior, understand which rule owns the conflict, and download a NAT analysis report for review.

NAT Policy Analyzer and Optimizer requirements

- NAT Policy Analyzer and Optimizer is available only when AIOps is enabled for the tenant and for the selected data source.
- For On-Premises Firewall Management Center, NAT Policy Analyzer and Optimizer requires Version 7.6 or later.

- If NAT Policy Analyzer and Optimizer is disabled in AIOps settings, NAT Policy Analyzer and Optimizer operations are blocked for all data sources.
- When AIOps is enabled for a data source, Policy Analyzer and Optimizer initiates the initial NAT synchronization and analysis for onboarded data sources.
- NAT Policy Analyzer and Optimizer performs analysis at the management center level, not on a per-device basis.



Note Disabled NAT rules are excluded from being analyzed by the Policy Analyzer and Optimizer.

Analyze NAT policies

Network Address Translation (NAT) Policy Analyzer and Optimizer analyzes Cloud-Delivered Firewall Management Center and On-Premises Firewall Management Center NAT policies and categorizes NAT findings into shadowed rules and redundant rules.

NAT Policy Analyzer and Optimizer reports these NAT finding types.

Table 149:

NAT finding	Meaning	How to interpret it
Shadowed NAT rule	A rule that will never evaluate network traffic because another rule that precedes it over shadows this rule.	Review the shadowing rule and the shadowed rules together. A rule is marked shadowed only when the overlap is complete for the compared dimensions.
Redundant NAT rule	A rule can be removed without changing the final NAT behavior because another rule can handle the traffic with the same effective NAT action.	Policy Analyzer and Optimizer checks the translated action for redundancy. Matching traffic criteria alone is not enough when translated source, destination, service, or PAT behavior differs.

NAT rule sections and comparison behavior

NAT Policy Analyzer and Optimizer follows the high-level NAT evaluation model used by Cisco Secure Firewall NAT policies. NAT rules are evaluated in section order: Manual NAT before Auto NAT, Auto NAT, and Manual NAT after Auto NAT.

The NAT section determines the rule comparisons that can be reported by NAT Policy Analyzer and Optimizer.

NAT section	What it contains	NAT Policy Analyzer and Optimizer comparison behavior
Manual NAT before Auto NAT	Manual or twice NAT rules evaluated before Auto NAT.	Rules in this section can be compared with rules in the same section, Auto NAT, and Manual NAT after Auto NAT.
Auto NAT	Object NAT rules whose device order is determined internally.	Auto NAT rules can participate in comparisons with Manual NAT before or after sections. NAT Policy Analyzer and Optimizer does not report shadowed or redundant rules within the same Auto NAT section.
Manual NAT after Auto NAT	Manual or twice NAT rules evaluated after Auto NAT.	Rules in this section can be compared with rules in the same section and with rules from earlier comparable sections.

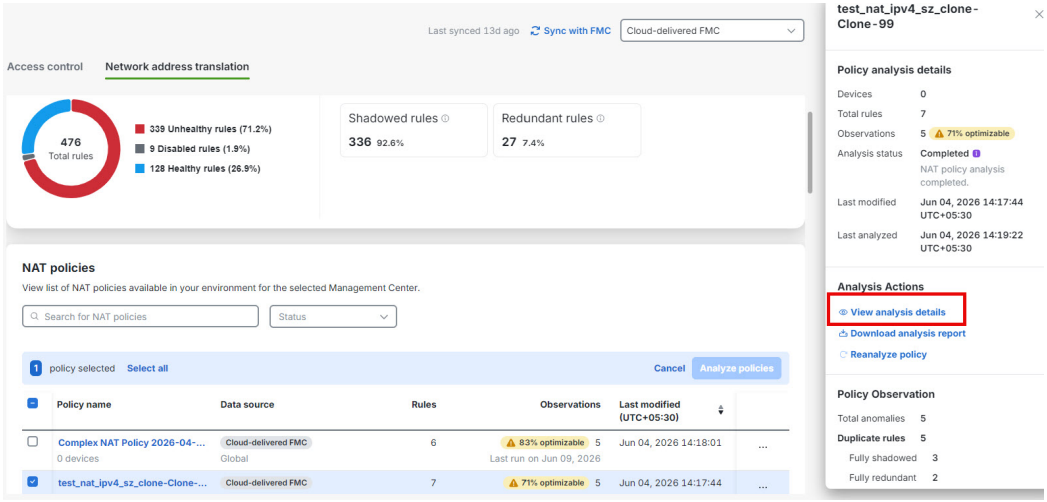
Supported NAT address families for Policy Analyzer and Optimizer

NAT Policy Analyzer and Optimizer supports these address-family flows.

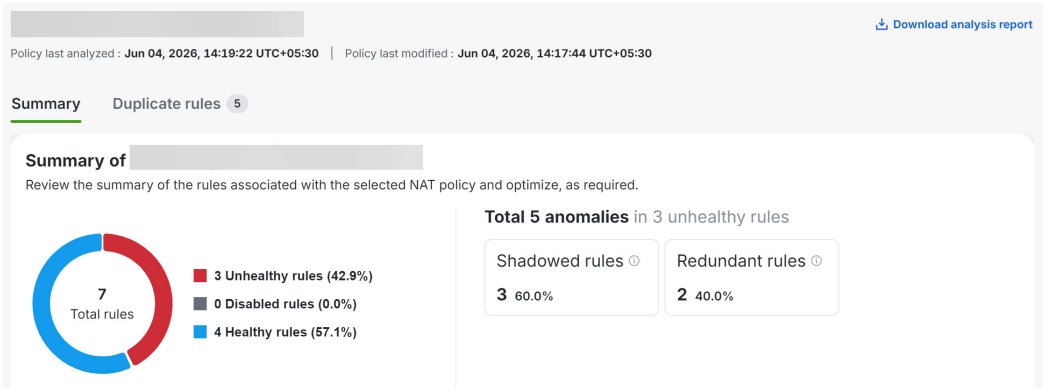
Flow	Traffic family	Notes
NAT44	IPv4 to IPv4	Supported for standard IPv4 NAT and PAT behavior.
NAT66	IPv6 to IPv6	Supported for IPv6 NAT behavior.
NAT64	IPv6 to IPv4	Supported when cross-family behavior is explicitly configured.
NAT46	IPv4 to IPv6	Supported when cross-family behavior is explicitly configured.

NAT policy analysis summary

When analysis completes, select a NAT policy to review the right-pane summary and click **View analysis details**.



The policy-specific page displaying the anomaly details is shown.



- **Overall summary**—provides insights on how many rules are unhealthy, disabled, healthy, using a pie chart for the selected NAT policy. You can also hover over the part of the pie to view the percentage of rules.
- **Total anomalies**—shows the number of rule anomalies detected in the selected NAT policy. Also indicates how many unhealthy rules contain those anomalies.

Download NAT analysis report

The NAT analysis report is available after NAT policy analysis is complete and can be downloaded as a PDF. The PDF report includes policy metadata, analysis timestamps, summary charts, counts for healthy, disabled, shadowed, and redundant rules, and tabular observation details.

1. In the Policy Analyzer and Optimizer page, click **Network address translation**.
2. Select a NAT policy, and in the right pane, click **Download analysis report**.

The detailed rule information can include NAT section, rule type, direction, source and destination interface objects, original source, original destination, original service, translated source, translated destination, and translated service. Use these details to evaluate whether changes are needed in the management center. NAT Policy Analyzer and Optimizer does not apply NAT remediations in this release.

NAT policy remediation

Policy Analyzer and Optimizer stages NAT remediation changes before it updates the policy. You can stage remediation for individual rules or for all rules in a duplicate-rule observation.

The available actions depend on the NAT rule type:

- Manual NAT rules: You can delete or disable selected manual rules
- Auto NAT rules: You can only delete Auto NAT rules, but you cannot disable them.

Note Back up your policy before applying any remediation. The changes made modify the NAT policy and cannot be automatically reverted.

Before you begin

- Ensure that NAT policy analysis is complete and that duplicate-rule observations are available.
- Verify the **Policy last analyzed** and **Policy last modified** values for the policy that you want to remediate.

Procedure

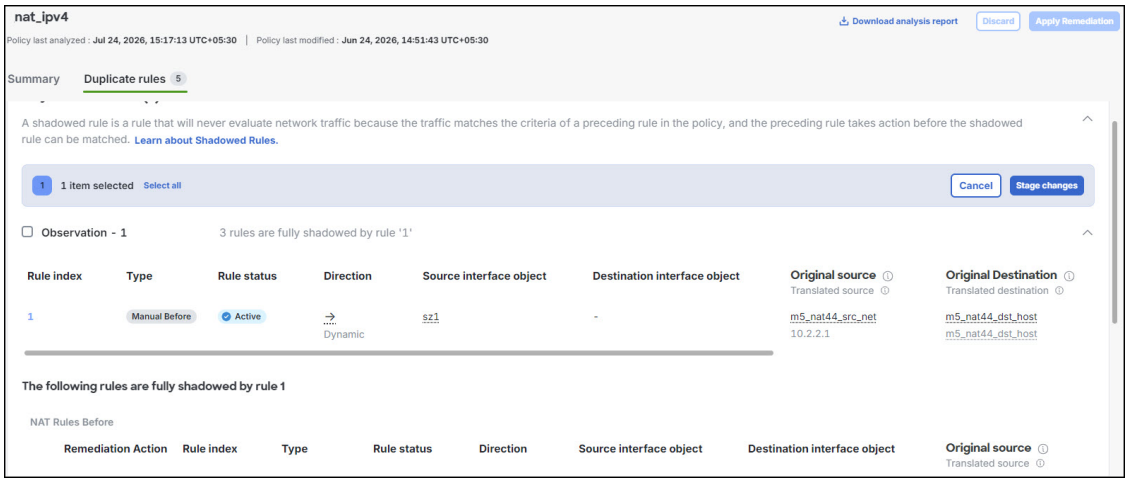
Step 1 In Policy Analyzer and Optimizer, select the NAT policy and under **Analysis Actions**, click **View analysis details** to open its analysis details.

Step 2 Click **Duplicate rules** and expand fully shadowed or redundant rules section that contains anomalies.

Step 3 Expand the observation that contains the rules that you want to remediate.

You can select the rules or select the observation to stage all rules in that observation.

Figure 396: NAT rule remediation



Step 4 Click **Stage changes**.

Step 5 In **Select rules to remediate**, select the action for each rule type:

NAT rule type	Available action
Manual rules	**Delete all selected** or **Disable all selected**
Auto NAT Rules	**Delete all selected**. You cannot disable **Auto NAT** rules.

Step 6 Select **Confirm staging**. The staged changes are only for review purposes and are not applied yet.

Note
Staged changes are visible to any user who opens the policy in Policy Analyzer and Optimizer. They remain until a user selects **Discard** or **Apply Remediation**, or until a new analysis runs after the policy is edited.

Click **View details** to inspect and review the staged changes. You can click **Undo** or **Discard** to revise the staged remediation.

Step 7 Select **Apply Remediation**.

Step 8 On the **Apply NAT duplicate-rule remediation** page, verify the policy and the number of rules to modify, and then select **Apply remediation**.

After you apply the remediation, Policy Analyzer and Optimizer updates the selected NAT rules and automatically starts a new analysis because the previous analysis is stale. While the new analysis runs, the **Observations** column for the NAT policy shows In progress. After analysis completes, review the refreshed duplicate-rule observations and policy status. If no analyzed anomalies remain, the policy can show a healthy status.

NAT remediation history

After you apply remediation to a NAT policy, Policy Analyzer and Optimizer records the remediation as a version in remediation history. Use remediation history to review when the remediation ran, the job that ran it, the user who started it, the remediation status, and the report for that version. The history lists the latest version first.



Note Remediation history records changes applied to the NAT policy. It does not indicate whether the updated policy has been deployed to managed devices. Deploy the policy separately as required.

Before you begin

Apply remediation to the NAT policy. Remediation history is available after the policy has at least one remediation version.

Procedure

Step 1 Open Policy Analyzer and Optimizer and click the Network address translation tab.

Step 2 Select the NAT policy that you want to review. The policy analysis details open.

Step 3 Under **Remediation Actions**, select **View remediation history**. The link shows the number of remediation versions that are available.

Step 4 Review the remediation history. The list displays the latest remediation version first and then the older versions.

Step 5 Review the date and time, Job ID, user, download option, and Remediation Status for each version.

For a completed remediation, select **Download** to download the remediation report.

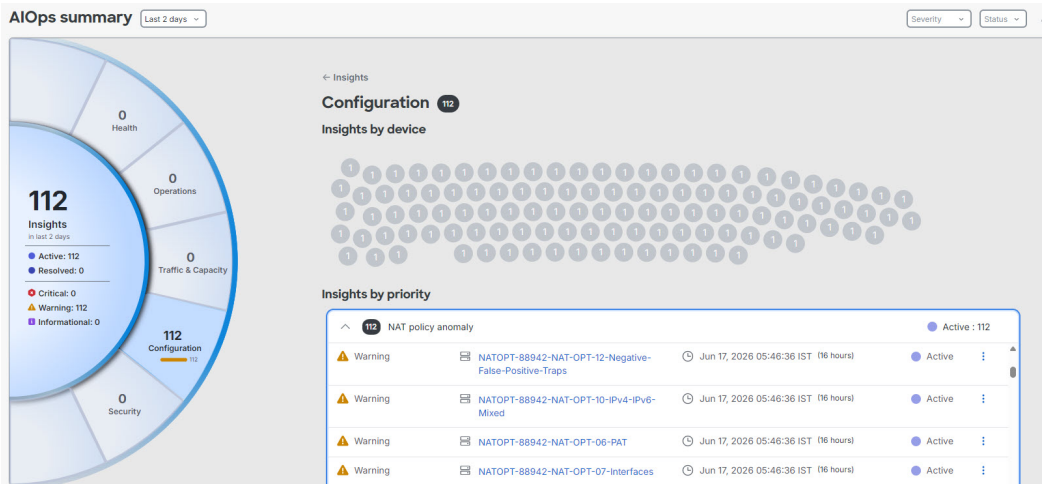
The remediation report is a PDF that summarizes the selected NAT remediation, its status, and the resulting changes to the policy.

Note
Policy Analyzer and Optimizer displays up to 10 remediation versions, from the latest version to the oldest available version.

NAT Policy Analyzer and Optimizer and AgenticOps insights

When NAT Policy Analyzer and Optimizer analysis detects NAT policy anomalies, AgenticOpscan raise a NAT policy anomaly insight. From AIOps Insights, you can review NAT anomaly counts by data source, open the affected data source or policy, and navigate back to the NAT Policy Analyzer and Optimizer analysis details.

Navigation: **Insights & Reports** > **Summary**.



If you dismiss a NAT anomaly insight, it is dismissed for the current insight occurrence. A later scheduled analysis can raise the insight again if NAT anomalies are detected again.

Troubleshooting Policy Analyzer and Optimizer

Read the following sections to troubleshoot any issues with the Policy Analyzer and Optimizer:

Policy Analyzer and Optimizer Does Not Analyze Policies

If you notice that Policy Analyzer and Optimizer is not analyzing policies despite clicking **Analyze Policy**, try the following:

Procedure

Step 1 Navigate **Administration** > **Integrations** > **Firewall Management Center**.

Step 2 Select the On-Premises Firewall Management Center or **Cloud-Delivered FMC** for which the policy analysis is not happening and choose **Workflows** under **Actions** on the right pane.

Step 3 If you see that the latest workflow's **Current State** shows up as **Error**, expand the workflow and scroll to the last action whose **END STATE** is **ERROR**.

Step 4 Click **Error Message** under the **RESULT** column to see a detailed error message or click **Stack Trace** to see the series of exceptions that occurred, which caused the error.

Step 5 Resolve the error or contact Cisco TAC for assistance.

Policy Analyzer and Optimizer Does Not Fetch Policies

If policies on your On-Premises Firewall Management Center are not displayed on the Policy Analyzer and Optimizer page on Security Cloud Control, do the following:

Procedure

Step 1 On the On-Premises Firewall Management Center, navigate **Integration** > **Cisco Security Cloud**.

Step 2 Ensure that the **Enable Policy Analyzer and Optimizer** checkbox is checked.

Step 3 (Optional) In the left navigation pane of your Security Cloud Control tenant, navigate **Administration** > **Integrations** > **Firewall Management Center**, and ensure that the On-Premises Firewall Management Center is active and reachable.

Frequently Asked Questions About Policy Analyzer and Optimizer

Can Cisco AI Assistant analyze and remediate policies instead of manually doing it using Policy Analyzer and Optimizer?

The Cisco AI Assistant collaborates with Policy Analyzer and Optimizer to scrutinize policies with anomalies and notify users. However, the AI Assistant cannot automatically analyze and remediate policies.

Can Policy Analyzer and Optimizer detect new changes to an already-analyzed policy and run analysis again on the same policy?

No, the Policy Analyzer and Optimizer can analyze policies only when manually triggered or at a 24-hour scheduled policy analysis run.

For a shared policy, does the Policy Analyzer and Optimizer provide individual device-based reports?

No. The Policy Analyzer and Optimizer provides reports only based on the access policy analysis data.

I am an On-Premises Firewall Management Center user. Should I purchase the Security Cloud Control base license to use the Policy Analyzer and Optimizer?

No. The Policy Analyzer and Optimizer comes as part of an existing or a newly created Security Cloud Control tenant during the Cisco Security Cloud integration.

I provisioned a Security Cloud Control tenant when I integrated my On-Premises Firewall Management Center with the Cisco Security Cloud. What other features, except Policy Analyzer and Optimizer, can I leverage in Security Cloud Control?

You can only leverage Policy Analyzer and Optimizer capabilities of this Security Cloud Control tenant. To use other features of Security Cloud Control, you need to purchase the Security Cloud Control base license and other device-specific licenses.

For an On-Premises Firewall Management Center on which the change management workflow is enabled and there are policies with pending changes to be approved, can the Policy Analyzer and Optimizer still apply remediations those policies?

No. The remediation will be hindered with an error saying the policies are locked for use.

Is there a maximum number of rules that Policy Analyzer and Optimizer can analyze in a policy?

There are no such limits. The Policy Analyzer and Optimizer can analyze any number of policies and rules. However, when the policies have more number of rules, the analysis takes a long time too.

What is the difference between disable rules and delete rules? Which is the better option?

Deleting a rule removes the rule completely from the device memory. However, disabling a rule keeps it in the device memory as a backup and does not get deployed to the device.

If a policy remediation fails when it is partially done, are the changes automatically revoked by Policy Analyzer and Optimizer?

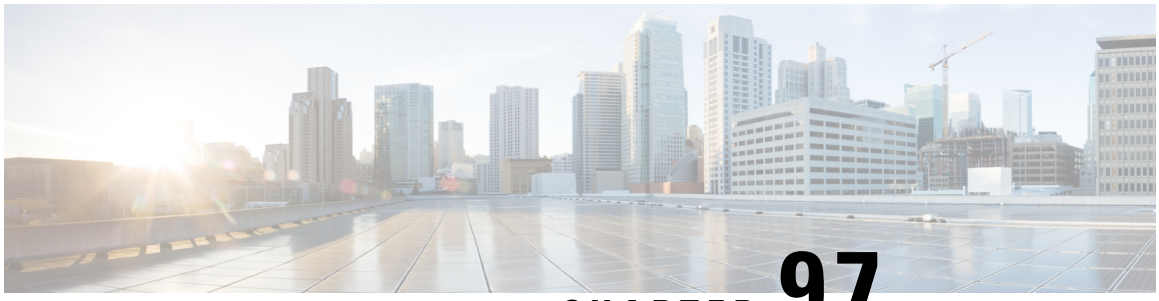
No. In such a case, you get a failure notification and a remediation report. You can read the report to know which rules were impacted by the half-done remediation, manually revoke the changes, and start the remediation all over again.



PART XXI

Custom Intrusion Policies for Access Control

- Intrusion Prevention , on page 2775
- Migrate from Snort 2 to Snort 3, on page 2799
- Intrusion Policies, on page 2819
- Intrusion Prevention Performance Tuning, on page 2865
- Network Analysis Policies, on page 2879
- Advanced Access Control settings for Network Analysis and Intrusion Policies, on page 2917



CHAPTER 97

Intrusion Prevention

The Snort inspection engine is an integral part of the Firewall Threat Defense. In most cases, the system-provided policies are all you need. For advanced users with specific needs, you can create custom rules. To configure custom Snort 3 rules, see Custom Snort 3 Intrusion Policies for Access Control.

Note To configure custom Snort 2 rules, see Custom Snort 2 Intrusion Policies for Access Control.

- Intrusion prevention, on page 2775
- Snort inspection engine, on page 2777
- Snort 3, on page 2777
- Multi-layer inspection, on page 2781
- Guidelines and limitations for network analysis and intrusion policies, on page 2784
- Prerequisites for network analysis and intrusion policies, on page 2784
- How policies examine traffic for intrusions, on page 2785
- System-provided and custom network analysis and intrusion policies, on page 2790

Intrusion prevention

Network analysis and intrusion policies work together as part of the intrusion detection and prevention feature.

- The term 'intrusion detection' generally refers to the process of passively monitoring and analyzing network traffic for potential intrusions and storing attack data for security analysis. This is sometimes referred to as "Intrusion Detection System (IDS)."
- The term 'intrusion prevention' includes the concept of intrusion detection, but adds the ability to block or alter malicious traffic as it travels across your network. This is sometimes referred to as "Intrusion Prevention System (IPS)."
- uses network analysis and intrusion policies working together to provide comprehensive packet inspection and threat protection.

Intrusion prevention components and operation

In an intrusion prevention deployment, when the system examines packets:

- A **network analysis policy** governs how traffic is *decoded* and *preprocessed* so it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.
- An **intrusion policy** uses *intrusion and preprocessor rules* (sometimes referred to collectively as *intrusion rules*) to examine the decoded packets for attacks based on patterns. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.

Both network analysis and intrusion policies are invoked by a parent access control policy, but at different times. As the system analyzes traffic, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (additional preprocessing and intrusion rules) phase. Together, network analysis and intrusion policies provide broad and deep packet inspection. They can help you detect, alert on, and protect against network traffic that could threaten the availability, integrity, and confidentiality of hosts and their data.

The system is delivered with several similarly named network analysis and intrusion policies (for example, Balanced Security and Connectivity) that complement and work with each other. By using system-provided policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and inspector rule states, as well as provides the initial configurations for inspectors and other advanced settings.

You can also create custom network analysis and intrusion policies. You can tune settings in custom policies to inspect traffic in the way that matters most to you so that you can improve both the performance of your managed devices and your ability to respond effectively to the events they generate.

You create, edit, save, and manage network analysis and intrusion policies using similar policy editors in the web interface. When you are editing either type of policy, a navigation panel appears on the left side of the web interface; the right side displays various configuration pages.

Refer to the videos for additional support and information:

- Snort 3 Condensed Overview
- Snort 3 Extended Overview

Attention Detection mode deprecation: From management center Version 7.4.0, for a network analysis policy (NAP), the Detection inspection mode is deprecated and will be removed in an upcoming release.

The Detection mode was intended to be used as a test mode so that you can enable inspections and see how they behave in your network before setting it to drop traffic, that is, to show traffic that would be dropped.

This behavior is improved where all inspector drops are controlled by the rule state, and you can set each one to generate events. This is done to test them before configuring the rule state to drop traffic. As we now have granular control over traffic drops in Snort 3, the Detection mode only adds more complexity to the product and is not needed, so the detection mode is deprecated.

If you change a NAP in Detection mode to Prevention, the NAP that processes the traffic of intrusion events and have the result "will be dropped" will now be "dropped" and the corresponding traffic will drop the traffic from these events. This is applicable for rules whose GIDs are not 1 or 3. GIDs 1 and 3 are text/compiled rules (typically provided by Talos or from your custom/imported rules) and all other GIDs are inspections for anomalies. These are more uncommon rules to trigger in a network. Changing to Prevention mode is unlikely to have any impact on the traffic. You need to just disable the intrusion rule that is applicable for the dropped traffic and set it to just generate or disable.

We recommend you choose Prevention as the inspection mode, but if you choose Prevention, you cannot revert to Detection mode.

Snort inspection engine

A Snort inspection engine is an integral traffic analysis component that

- analyzes traffic in real time to provide deep packet inspection,
- operates as part of the Secure Firewall Threat Defense (formerly Firepower Threat Defense) device, and
- enables network analysis and intrusion policies to detect and protect against intrusions.

Snort 3

Snort 3 is a network inspection engine that

- provides vast improvements compared to the earlier version of Snort,
- offers more efficient operation with better performance and scalability, and
- is architecturally redesigned to inspect more traffic with equivalent resources when compared to Snort 2.

Snort 3 changes and improvements

The older version of Snort is Snort 2. Snort 3 provides simplified and flexible insertion of traffic parsers. Snort 3 also provides new rule syntax that makes rule writing easier and shared object rule equivalents visible.

The other significant changes with Snort 3 are:

- Unlike Snort 2, which uses multiple Snort instances, Snort 3 associates multiple threads with a single Snort instance. This uses less memory, improves Snort reload times, and supports more intrusion rules and a larger network map. The number of Snort threads varies by platform and is the same as the number of Snort 2 instances for each platform. Usage is virtually transparent.
- Snort version per Firewall Threat Defense—The Snort inspection engine is Firewall Threat Defense specific and not Cloud-Delivered Firewall Management Center (formerly Firepower Management Center) specific. Cloud-Delivered Firewall Management Center can manage several Firewall Threat Defenses, each with either versions of Snort (Snort 2 and Snort 3). Although the Cloud-Delivered Firewall Management Center's intrusion policies are unique, the system applies Snort 2 or Snort 3 version of an intrusion policy for intrusion protection depending on the device's selected inspection engine.
- Decoder rules—Packet decoder rules fire only in the default intrusion policy. The system ignores decoder rules that you enable in other policies.
- Shared object rules—Snort 3 supports some, but not all, shared object (SO) intrusion rules (rules with a generator ID (GID) of 3). Enabled shared object rules that are not supported do not trigger.
- Multi-layer inspection for Security Intelligence—Snort 3 detects the innermost IP address regardless of the layer.
- Platform support—Snort 3 requires Firewall Threat Defense version 7.0 or later. It is not supported with ASA FirePOWER or NGIPSv.
- Managed Devices—An Cloud-Delivered Firewall Management Center with version 7.0 can simultaneously support versions 6.4, 6.5, 6.6, 6.7, and 7.0 Snort 2 Firewall Threat Defenses, and version 7.0 Snort 3 Firewall Threat Defenses.
- Traffic interruption when switching Snort versions—Switching Snort versions interrupts traffic inspection and a few packets might drop during deployment.
- Unified policies—Irrespective of the underlying Snort engine version that is enabled in the managed Firewall Threat Defenses, the access control policies, intrusion policies, and network analysis policies configured in the Cloud-Delivered Firewall Management Center work seamlessly in applying the policies. All intrusion policies in Cloud-Delivered Firewall Management Center version 7.0 and above have two versions available, Snort 2 version and Snort 3 version. The intrusion policy is unified, which means that it has a common name, base policy, and inspection mode, although there are two versions of the policy (Snort 2 version and Snort 3 version). The Snort 2 and the Snort 3 versions of the intrusion policy can be different in terms of the rule settings. However, when the intrusion policy is applied on a device, the system automatically identifies the Snort version enabled on the device and applies the rule settings configured for that version.
- Lightweight Security Package (LSP)—Replaces the Snort Rule Updates (SRU) for Snort 3 next-generation intrusion rule and configuration updates. Downloading updates downloads both the Snort 3 LSP and the Snort 2 SRU.

  LSP updates provide new and updated intrusion rules and inspector rules, modified states for existing rules, and modified default intrusion policy settings for Cloud-Delivered Firewall Management Center and Firewall Threat Defense versions 7.0 or above. When you upgrade an Cloud-Delivered Firewall Management Center from version 6.7 or lower to 7.0, it supports both LSPs and SRUs. LSP updates may also delete system-provided rules, provide new rule categories and default variables, and modify default variable values. For more information on LSP updates, see the *Update Intrusion Rules* topic in the latest version of the *Firepower Management Center Configuration Guide*.

- Mapping of Snort 2 and Snort 3 rules and presets—Snort 2 and Snort 3 rules are mapped and the mapping is system-provided. However, it is not a one-to-one mapping. The system-provided intrusion base policies are pre-configured for both Snort 2 and Snort 3, and they provide the same intrusion prevention although with different rule sets. The system-provided base policies for Snort 2 and Snort 3 are mapped with each other for the same intrusion prevention settings.
- Synchronizing Snort 2 and Snort 3 rule override—When a Firewall Threat Defense is upgraded to 7.0, you can upgrade the inspection engine of the Firewall Threat Defense to the Snort 3 version. Cloud-Delivered Firewall Management Center maps all the overrides in the existing rules of the Snort 2 version of the intrusion policies to the corresponding Snort 3 rules using the mapping provided by Talos. However, if there are additional overrides performed after the upgrade or if you have installed a new Firewall Threat Defense of version 7.0, they have to be manually synchronized.
- Custom intrusion rules—You can create custom intrusion rules in Snort 3. You can also import the custom intrusion rules that exist for Snort 2 to Snort 3.
- Rule groups—The Cloud-Delivered Firewall Management Center groups all Snort 3 rules into rule groups. Rule groups are logical groups of rules that provide an easy management interface to enhance rule accessibility, rule navigation, and better control over the rule group security level.

  From Cloud-Delivered Firewall Management Center 7.3.0, rule navigation for several levels of rule groups is supported that provides more flexibility and logical grouping of rules. The MITRE framework is added that enables you to navigate through rules using the MITRE framework. MITRE is just another category of rule groups and is a part of Talos rule groups.

Note See https://attack.mitre.org for information about MITRE.

  A rule can be part of multiple rule groups, such as multiple MITRE ATT&CK rule groups, a rule category rule group, multiple "asset type" rule groups, a malware campaign, and others. The available rule groups are listed in the intrusion policy editor and can be selected to enhance policies.

  With this multi-level hierarchical structure, you can traverse down to the last element, which is the "leaf rule group." These rule groups contain sets of rules that are related to each other, such as a specific type of vulnerability, a similar target system, or a similar threat category. Rule groups have four security levels associated with them. You can change the security level, add or remove rule groups, and you can change the rule action for rules that match traffic seen on the network. This is done to bring a satisfactory balance between security, performance, and false positive resistance.
- Switching between Snort 2 and Snort 3 engines—Firewall Threat Defenses that support Snort 3 can also support Snort 2. Switching from Snort 3 to Snort 2 is not recommended from the efficacy perspective.

Important Although you can switch Snort versions freely, intrusion rule changes in one version of Snort will not be updated in the other version automatically. If you change the rule action for a rule in one version of Snort, ensure you replicate the change in the other version before switching the Snort version. System provided synchronization option only synchronizes the changes in the Snort 2 version of the intrusion policy to the Snort 3 version, and not the other way around.

Snort 3 multi-process support

Snort 3 multi-process support is a feature that:

- associates multiple threads with two Snort instances on Secure Firewall 6160 and 6170 devices running Secure Firewall version 10.0,
- automatically configures each Snort 3 instance with the threads as required during deployment, and
- results in reduced memory load, reduced lock contention, and shorter core generation times for improved resiliency, performance, and scalability.

Snort 3 multi-process architecture

Individual Snort 3 instances provide several advantages. Any Snort instance failure results in minimal traffic impact and you can also restart a single Snort process, if required.

Snort threads are distributed over NUMA nodes where each node is considered as an independent processor. This ensures that memory boundaries are maintained. For example, Snort instance 0 uses the memory of node 0, and instance 1 uses the memory of node 1.

The CLI outputs of commands such as **show coredump** and **show perfstats** have been enhanced to display information on both the Snort instances. On the Management Center, navigate to **Troubleshooting > Monitor > Devices** to view the **Overview > Critical Processes** and the **Memory** sections for information on both the Snort instances. If there are any rule profiling errors, the Snort instance number is displayed with the error.

Disable snort 3 multi-process support

This task disables Snort 3 multi-process support.

Disabling multi-process support is intended for debugging purposes only and requires a manual deployment to take effect.

Procedure

Step 1 Use the **configure snort multi-process disable** command to disable Snort 3 multi-process support.

Example:

```
> configure snort multi-process disable
----------------------------------------------------------------------
Caution: this command is intended for debugging purposes only.
A deployment is required after enabling/disabling multi-process and will cause a Snort
restart and temporary traffic interruption.
----------------------------------------------------------------------
Do you still want to continue?
Please enter 'YES' or 'NO': YES
Please perform a manual deployment to disable multi-process.
```

Note
This command causes the Snort engine to restart and will also lead to traffic interruption during the next deployment when Snort is changing from the multi-process state to a single-process state.

Step 2 Redeploy the configuration to ensure that the Snort 3 multi-process support is disabled.

Multi-layer inspection

The multi-layer inspection feature in Snort 3 enables you to inspect traffic across multiple protocol layers within the same session.

Snort 3 can simultaneously inspect the following layers:

- Network layer (L3): IP addresses, fragmentation
- Transport layer (L4): TCP or UDP ports, flags, session tracking
- Application layer (L7): HTTP, DNS, SMTP, TLS, SIP, and others.
- File and payload layer: File extraction, malware signatures, embedded content

Snort 3 supports decoding for a wide range of protocols. It inspects encapsulated protocols and evaluates packet headers to determine policy actions during advanced threat inspection.

Header inspection

On versions earlier than 7.6.4

For encapsulation protocols not supported by Lina, the system initially matches traffic based on the outer IP header. The subsequent evaluation depends on the rule criteria:

- Rules without Layer 7 (L7) criteria: If the matched rule does not include L7 criteria (such as URL, application ID, or user identity), the system applies the rule based on the outer IP header.
- Rules with Layer 7 (L7) criteria: If the rule includes L7 criteria, the system uses the outer-IP header match as a starting point and then performs an additional evaluation using the innermost header.

This behavior can lead to unexpected results, such as a non-L7 rule matching the outer-IP header taking precedence over a rule matching the inner-IP header located lower in the rule set.

On version 7.6.4 and later

- Snort 3 always inspects the innermost header for protocols not supported by Lina, regardless of the presence of Layer 7 rules.
- Outer header-based rules alone are not sufficient for policy enforcement for protocols such as VXLAN. Snort 3 disregards connection details for action in such cases, using the actual packet content instead.
- For protocols supported by Lina, policy action continues to be based on the innermost header.
- If you currently manage encapsulated traffic for protocols not supported by Lina (such as VXLAN) using only outer-IP rules, be aware that on version 7.6.4 or later, this traffic will be handled by the default access control rule. In such a scenario, if your default policy action is set to Block, this may result in the traffic being dropped. We recommend that you review and update access control policies to match the inner IP header for such protocols, in addition to any existing outer IP header configurations.

Layer 7 features

For encapsulated traffic in Layer 7 features, only the innermost IP header must be matched in rules. Outer header-only rules will not apply.

Example scenarios

The following table lists examples based on different parameters.

Protocol	Version	Outer IP rule	Inner IP rule	L7 rule needed	Result
VXLAN (not supported by LINA)	7.6.2	Yes	No	No	Traffic allowed if outer header matches
VXLAN (not supported by LINA)	7.6.4	No	Yes	No	Traffic allowed if inner header matches
GRE (supported by LINA)	7.6.2 and 7.6.4	No	Yes	No	Traffic allowed if inner header matches
VXLAN + L7 feature	7.6.2 and 7.6.4	No	Yes	Yes	Inner header rule required



Note An outer IP rule matches traffic based on the tunnel endpoints. It evaluates:

- Source IP address of the encapsulating device
- Destination IP address of the encapsulating device

An inner IP rule matches traffic based on the original IP packet inside the tunnel. It evaluates:

- The real client IP address
- The real destination server IP address

Encapsulation layer limits and inspection behavior

When managing encapsulated traffic, it is important to understand the processing limits of the system's two primary inspection engines, Lina and Snort.

- LINA: Supports decoding up to 2 layers of encapsulation.
- Snort: Supports decoding up to 255 layers of encapsulation (with a default limit of 40).

Supported protocols and inspection capabilities of LINA and Snort 3

LINA

LINA supports decoding for the following protocols across the Data Link (L2), Network (L3), and Transport (L4) layers.

Layer	Supported protocols and headers
Layer 2	Cisco MetaData (CMD), IEEE 802.1Q, IEEE 802.1ad (DOT1AD), and Q-in-Q (EtherTypes 0x9100 and 0x9200). Supports up to three levels of VLAN headers.
Layer 3	IPv4, IPv6
Layer 4	ICMP (v4/v6), ESP, VXLAN, Geneve, GRE, IP-in-IP, and IPv6-in-IPv6. Supports up to 1 level of encapsulation.

GRE-specific support

For GRE traffic, the system supports decoding for the following protocol types:

- Point-to-Point Tunneling Protocol (PPTP)
- Web Cache Communication Protocol (WCCP)

ERSPAN handling

The system decodes GRE and ERSPAN (types 1, 2, and 3). To facilitate dispatch layer load-balancing, the system internally skips to the inner headers. Consequently, other modules are unable to inspect ERSPAN headers.

Snort 3

While Snort 3 supports a vast array of protocols, the following list highlights common encapsulation and tunneling protocols supported for inspection:

- VXLAN
- GRE
- Teredo
- GTP
- MPLS
- Geneve
- IP-in-IP
    - IPv4-in-IPv4
    - IPv6-in-IPv4
    - IPv4-in-IPv6
    - IPv6-in-IPv6

Note This list is not exhaustive. Snort 3 is designed to decode a wide range of industry-standard protocols.

Guidelines and limitations for network analysis and intrusion policies

Feature Limitations of Snort 3 for Firewall Management Center-Managed Firewall Threat Defense

The table lists the features that are supported on Snort 2 but not supported on Snort 3 for Firewall Management Center-managed Firewall Threat Defense devices.

Table 150: Feature Limitations of Snort 3

Policy/Area	Features not supported
Access Control Policy	The following application settings: • Safe Search • YouTube EDU
Intrusion Policy	• Global rule thresholding • Logging configuration: • SNMP • SRU rule updates as Snort 3 supports only LSP rule updates
Other features	Event logging with fully qualified domain name (FQDN).

Prerequisites for network analysis and intrusion policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the Intrusion Prevention System (IPS) license enabled for the Firewall Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

How policies examine traffic for intrusions

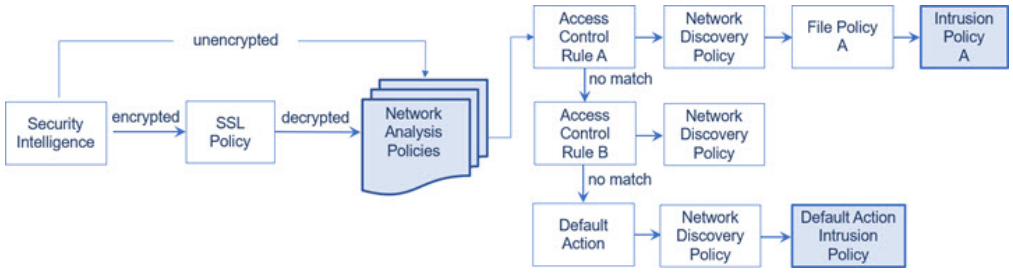
When the system analyzes traffic as part of your access control deployment, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (intrusion rules and advanced settings) phase.

Summary

The key components involved in examining traffic for intrusions are:

- Security Intelligence: Performs initial traffic filtering and blocking
- SSL policy: Handles encrypted traffic inspection and decryption
- Network analysis policies: Conduct decoding and preprocessing of traffic
- Access control rules: Determine which policies to apply to traffic
- Intrusion policies: Examine traffic for security threats and intrusions
- File policies: Analyze files for malware and threats
- Network discovery policy: Passively inspects packets for network mapping

Workflow



These stages describe how the system examines traffic for intrusions in an inline deployment:

1. The system processes traffic through multiple policy layers in a specific order, with network analysis and intrusion policy selection phases highlighted in the deployment.

   In an inline deployment (that is, where relevant configurations are deployed to devices using routed, switched, or transparent interfaces, or inline interface pairs), the system can block traffic without further inspection at almost any step in the illustrated process. Security Intelligence, the SSL policy, network analysis policies, file policies, and intrusion policies can all either drop or modify traffic. Only the network discovery policy, which passively inspects packets, cannot affect the flow of traffic.

2. At each step of the process, a packet could cause the system to generate an event. Intrusion and preprocessor events (sometimes referred to collectively as *intrusion events*) are indications that a packet or its contents may represent a security risk.

Tip The diagram does not reflect that access control rules handle encrypted traffic when your SSL inspection configuration allows it to pass, or if you do not configure SSL inspection. By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

3. Note that for a single connection, although the system selects a network analysis policy before an access control rule as shown in the diagram, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

Decoding, normalizing, and preprocessing: Network analysis policies

Without decoding and preprocessing, the system cannot appropriately evaluate traffic for intrusions because protocol differences would make pattern matching impossible. Network analysis policies govern these traffic-handling tasks

- after traffic is filtered by Security Intelligence,
- after encrypted traffic is decrypted by an optional SSL policy, and
- before traffic can be inspected by file or intrusion policies

A network analysis policy governs packet processing in phases. First the system decodes packets through the first three TCP/IP layers, then continues with normalizing, preprocessing, and detecting protocol anomalies:

- The packet decoder converts packet headers and payloads into a format that can be easily used by the inspectors and later, intrusion rules. Each layer of the TCP/IP stack is decoded in turn, beginning with the data link layer and continuing through the network and transport layers. The packet decoder also detects various anomalous behaviors in packet headers.
- In inline deployments, the inline normalization preprocessor reformats (normalizes) traffic to minimize the chances of attackers evading detection. It prepares packets for examination by other inspectors and intrusion rules, and helps ensure that the packets the system processes are the same as the packets received by the hosts on your network.
- Various network and transport layers inspectors detect attacks that exploit IP fragmentation, perform checksum validation, and perform TCP and UDP session preprocessing.

  Note that some advanced transport and network inspector settings apply globally to all traffic handled by the target devices of an access control policy. You configure these in the access control policy rather than in a network analysis policy.
- Various application-layer protocol decoders normalize specific types of packet data into formats that the intrusion rules engine can analyze. Normalizing application-layer protocol encodings allows the system to effectively apply the same content-related intrusion rules to packets whose data is represented differently, and to obtain meaningful results.
- The Modbus, DNP3, CIP, and s7commplus SCADA inspectors detect traffic anomalies and provide data to intrusion rules. Supervisory Control and Data Acquisition (SCADA) protocols monitor, control, and acquire data from industrial, infrastructure, and facility processes such as manufacturing, production, water treatment, electric power distribution, airport, and shipping systems.

- Several inspectors allow you to detect specific threats, such as Back Orifice, portscans, SYN floods and other rate-based attacks.

  Note that you configure the sensitive data inspector, which detects sensitive data such as credit card numbers and Social Security numbers in ASCII text, in intrusion policies.



Note When TLS Server Identity is disabled, Snort 3 does not perform SNI mismatch detection. It only evaluates the SNI in the Client Hello packet and bypasses the validation of the certificate's Common Name (CN) in the Server Hello packet.

In a newly created access control policy, one default network analysis policy governs preprocessing for all traffic for all intrusion policies invoked by the same parent access control policy. Initially, the system uses the Balanced Security and Connectivity network analysis policy as the default, but you can change it to another system-provided or custom network analysis policy. In a more complex deployment, advanced users can tailor traffic preprocessing options to specific security zones, networks, and VLANs by assigning different custom network analysis policies to preprocess matching traffic.



Note For an access control policy with rule action as Trust and a prefilter rule with action as Fastpath with logging options disabled, you will observe that the end-of-flow events are still generated in the system. The events are not visible on the management center event pages.

Access control rules: Intrusion policy selection

Intrusion policy selection is a traffic evaluation process that—

- evaluates traffic after initial preprocessing using access control rules,
- applies matching access control rule to handle traffic through monitor, trust, block, or allow actions,
- inspects allowed traffic for discovery data, malware, prohibited files, and intrusions in that order, and
- handles non-matching traffic using the access control policy's default action.

In most cases, the first access control rule that a packet matches is the rule that handles that traffic; you can monitor, trust, block, or allow matching traffic.

Traffic evaluation process

When you allow traffic with an access control rule, the system can inspect the traffic for discovery data, malware, prohibited files, and intrusions, in that order. Traffic not matching any access control rule is handled by the access control policy's default action, which can also inspect for discovery data and intrusions.



Note All packets, **regardless** of which network analysis policy preprocesses them, are matched to configured access control rules, making them potentially subject to inspection by intrusion policies in top-down order.

The diagram in How policies examine traffic for intrusions, on page 2785 shows the flow of traffic through a device in an inline, intrusion prevention and AMP for Networks deployment—

- Access Control Rule A allows matching traffic to proceed. The traffic is then inspected for discovery data by the network discovery policy, for prohibited files and malware by File Policy A, and then for intrusions by Intrusion Policy A.
- Access Control Rule B also allows matching traffic. However, in this scenario, the traffic is not inspected for intrusions (or files or malware), so there are no intrusion or file policies associated with the rule. Note that by default, traffic that you allow to proceed is inspected by the network discovery policy; you do not need to configure this.
- In this scenario, the access control policy's default action allows matching traffic. The traffic is then inspected by the network discovery policy, and then by an intrusion policy. You can (but do not have to) use a different intrusion policy when you associate intrusion policies with access control rules or the default action.

The example in the diagram does not include any blocking or trusting rules because the system does not inspect blocked or trusted traffic.

Intrusion inspection: Intrusion policies, rules, and variable sets

Intrusion inspection is a security mechanism that

- uses intrusion policies to govern how the system inspects traffic for security violations and can block or alter malicious traffic in inline deployments
- employs intrusion and preprocessor rules to analyze network traffic against specified criteria and detect attempts to exploit vulnerabilities, and
- utilizes variable sets to define values commonly used in intrusion rules for identifying source and destination IP addresses and ports.

Intrusion and inspector rules

An intrusion rule is a specified set of keywords and arguments that detects attempts to exploit vulnerabilities on your network; the system uses an intrusion rule to analyze network traffic to check if it matches the criteria in the rule. The system compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in a rule, the rule triggers.

The system includes these types of rules created by Cisco Talos Intelligence Group (Talos):

- *shared object intrusion rules*, which are compiled and cannot be modified, except for rule header information such as source and destination ports and IP addresses
- *standard text intrusion rules*, which can be saved and modified as new custom instances of the rule.
- *preprocessor rules*, which are rules associated with inspectors and packet decoder detection options in the network analysis policy. You cannot copy or edit inspector rules. Most inspector rules are disabled by default; you must enable them to use inspectors to generate events and, in an inline deployment, drop offending packets.

When the system processes packets according to an intrusion policy, first a rule optimizer classifies all activated rules in subsets based on criteria such as: transport layer, application protocol, direction to or from the protected network, and so on. Then, the intrusion rules engine selects the appropriate rule subsets to apply to each packet.

Finally, a multi-rule search engine performs three different types of searches to determine if the traffic matches the rule:

- The protocol field search looks for matches in particular fields in an application protocol.
- The generic content search looks for ASCII or binary byte matches in the packet payload.
- The packet anomaly search looks for packet headers and payloads that, rather than containing specific content, violate well-established protocols.

In a custom intrusion policy, you can tune detection by enabling and disabling rules, as well as by writing and adding your own standard text rules. You can also use Cisco recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.

Note When there are insufficient packets to process specific traffic against a block rule, the system continues to evaluate the remaining traffic against other rules. If any of the remaining traffic matches a rule which is set to block, then the session is blocked. However, if the system analyses the remaining traffic to be passed, then the traffic status shows pending on the rule which is stuck for want of complete packets.

Variable Sets

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Most variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

The system provides a single default variable set, which is comprised of predefined default variables. Most system-provided shared object rules and standard text rules use these predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.

Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends that you modify key default variables in the default set. When you use variables that accurately reflect your network environment, processing is optimized and the system can monitor relevant systems for suspicious activity. Advanced users can create and use custom variable sets for pairing with one or more custom intrusion policies.

Important If you are creating a custom variable set, do not use a number as the first character in a custom variable set name (for example, 3Snort). This will cause Snort 3 validation to fail when you deploy a configuration to Firewall Threat Defense firewall on the Cloud-Delivered Firewall Management Center.

Intrusion event generation

Intrusion event generation is a security mechanism that:

- identifies possible intrusions through packet decoding, preprocessing, and rules engine analysis,
- generates *intrusion or preprocessor events* (sometimes collectively called *intrusion events*) when threats are detected, and
- transmits events from managed devices to the Firewall Management Center for aggregated analysis and network asset protection.

Event composition and data collection

Each intrusion event in the database includes an event header and contains information about the event name and classification; the source and destination IP addresses; ports; the process that generated the event; and the date and time of the event, as well as contextual information about the source of the attack and the target. For packet-based events, the system also logs a copy of the decoded packet header and payload for the packet or packets that triggered the event.

As the database accumulates intrusion events, you can begin your analysis of potential attacks. The system provides you with the tools you need to review intrusion events and evaluate whether they are important in the context of your network environment and your security policies.

Event generation scenarios

The packet decoder, the preprocessors, and the intrusion rules engine can all cause the system to generate an event. For example:

- If the packet decoder (configured in the network analysis policy) receives an IP packet that is less than 20 bytes, which is the size of an IP datagram without any options or payload, the decoder interprets this as anomalous traffic. If, later, the accompanying decoder rule in the intrusion policy that examines the packet is enabled, the system generates a inspector event.
- If the IP defragmentation preprocessor encounters a series of overlapping IP fragments, the inspector interprets this as a possible attack and, when the accompanying inspector rule is enabled, the system generates a inspector event.
- Within the intrusion rules engine, most standard text rules and shared object rules are written so that they generate intrusion events when triggered by packets.

System-provided and custom network analysis and intrusion policies

Network analysis and intrusion policies are security mechanisms that

- examine traffic through preprocessing and intrusion prevention phases,
- provide system-provided default configurations for immediate deployment, and
- allow custom policy creation for tailored security requirements.

Default policy behavior in access control deployments

By default, a newly created access control policy invokes system-provided network analysis and intrusion policies to examine traffic.

This diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic:



The traffic flow demonstrates these characteristics:

- A default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided *Balanced Security and Connectivity network analysis policy* is the default.
- The default action of the access control policy allows all non-malicious traffic, as determined by the system-provided *Balanced Security and Connectivity intrusion policy*. Because the default action allows traffic to pass, the discovery feature can examine it for host, application, and user data before the intrusion policy can examine and potentially block malicious traffic.
- The policy uses default Security Intelligence options (global Block and Do Not Block lists only), does not decrypt encrypted traffic with an SSL policy, and does not perform special handling and inspection of network traffic using access control rules.

A simple step you can take to tune your intrusion prevention deployment is to use a different set of system-provided network analysis and intrusion policies as your defaults. Cisco delivers several pairs of these policies with the system.

Or, you can tailor your intrusion prevention deployment by creating and using custom policies. You may find that the inspector options, intrusion rule, and other advanced settings configured in those policies do not address the security needs of your network. By tuning your network analysis and intrusion policies, you can control at a very granular level how the system processes and inspects traffic on your network for intrusions.

System-provided network analysis and intrusion policies

System-provided network analysis and intrusion policies are security configurations that

- leverage the experience of the Cisco Talos Intelligence Group (Talos) for intrusion and inspector rule states and advanced settings,
- provide starting points for well tuned defensive policies covering common network setups, and
- receive automatic updates through Lightweight Security Package (LSP) releases for new vulnerabilities.

Key characteristics and usage

No system-provided policy covers every network profile, traffic mix, or defensive posture. Each covers common cases and network setups that provide a starting point for a well-tuned defensive policy. Although you can use system-provided policies as-is, Cisco strongly recommends that you use them as the base for custom policies that you tune to suit your network.

Tip Even if you use system-provided network analysis and intrusion policies, you should configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify key default variables in the default set.

As new vulnerabilities become known, Talos releases intrusion rule updates, also known as *Lightweight Security Package* (LSP). These rule updates can modify any system-provided network analysis or intrusion policy, and can provide new and updated intrusion rules and inspector rules, modified states for existing rules, and modified default policy settings. Rule updates may also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

If a rule update affects your deployment, the web interface marks affected intrusion and network analysis policies as out of date, as well as their parent access control policies. You must re-deploy an updated policy for its changes to take effect.

For your convenience, you can configure rule updates to automatically re-deploy affected intrusion policies, either alone or in combination with affected access control policies. This allows you to easily and automatically keep your deployment up-to-date to protect against recently discovered exploits and intrusions.

To ensure up-to-date preprocessing settings, you **must** re-deploy access control policies, which also deploys any associated SSL, network analysis, and file policies that are different from those currently running, and can also update default values for advanced preprocessing and performance options.

Cisco delivers the system-provided network analysis and intrusion policies:

- Balanced Security and Connectivity network analysis and intrusion policies — These policies are built for both speed and detection. Used together, they serve as a good starting point for most organizations and deployment types. The system uses the Balanced Security and Connectivity policies and settings as defaults in most cases.
- Connectivity Over Security network analysis and intrusion policies — These policies are built for organizations where connectivity (being able to get to all resources) takes precedence over network infrastructure security. The intrusion policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled.
- Security Over Connectivity network analysis and intrusion policies — These policies are built for organizations where network infrastructure security takes precedence over user convenience. The intrusion policy enables numerous network anomaly intrusion rules that could alert on or drop legitimate traffic.
- Maximum Detection network analysis and intrusion policies — These policies are built for organizations where network infrastructure security is given even more emphasis than is given by the Security Over Connectivity policies, with the potential for even greater operational impact. For example, the intrusion policy enables rules in a large number of threat categories including malware, exploit kit, old and common vulnerabilities, and known in-the-wild exploits.
- No Rules Active intrusion policy — In the No Rules Active intrusion policy, all intrusion rules, and all advanced settings except intrusion rule thresholds, are disabled. This policy provides a starting point if you want to create your own intrusion policy instead of basing it on the enabled rules in one of the other system-provided policies.

Note Depending on the system-provided base policy that is selected, the settings of the policy vary. To view the policy settings, click the **Edit** icon next to the policy and then click the **Base Policy** drop-down menu.

Benefits of custom network analysis and intrusion policies

Custom network analysis and intrusion policies are security configurations that:

- improve system performance in your specific environment,
- provide a focused view of malicious traffic and policy violations occurring on your network, and
- allow granular configuration of how the system processes and inspects network traffic for intrusions.

Custom policy structure

You may find that the inspector options, intrusion rules, and other advanced settings configured in the system-provided network analysis and intrusion policies do not fully address the security needs of your organization.

All custom policies have a base policy, also called a base layer, which defines the default settings for all configurations in the policy. A layer is a building block that you can use to efficiently manage multiple network analysis or intrusion policies.

In most cases, you base custom policies on system-provided policies, but you can use another custom policy. However, all custom policies have a system-provided policy as the eventual base in a policy chain. Because rule updates can modify system-provided policies, importing a rule update may affect you even if you are using a custom policy as your base. If a rule update affects your deployment, the web interface marks affected policies as out of date.

Custom network analysis policy benefits

Custom network analysis policies:

- enable tuning preprocessing to match specific traffic characteristics and monitoring requirements,
- reduce system processing overhead by disabling unnecessary inspectors, and
- allow tailored preprocessing for different security zones, networks, or VLANs in complex deployments.

By default, one network analysis policy preprocesses all unencrypted traffic handled by the access control policy. That means that all packets are decoded and preprocessed according to the same settings, regardless of the intrusion policy (and therefore intrusion rule set) that later examines them.

Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default. A simple way to tune preprocessing is to create and use a custom network analysis policy as the default.

Tuning options available vary by inspector, but some of the ways you can tune inspectors and decoders include:

- You can disable inspectors that do not apply to the traffic you are monitoring. For example, the HTTP Inspect inspector normalizes HTTP traffic. If you are confident that your network does not include any web servers using Microsoft Internet Information Services (IIS), you can disable the inspector option that looks for IIS-specific traffic and thereby reduce system processing overhead.

Note If you disable an inspector in a custom network analysis policy, but the system needs to use that inspector to later evaluate packets against an enabled intrusion or inspector rule, the system automatically enables and uses the inspector although the inspector remains disabled in the network analysis policy web interface.

- Specify ports, where appropriate, to focus the activity of certain inspectors. For example, you can identify additional ports to monitor for DNS server responses, encrypted SSL sessions, or ports on which you decode telnet, HTTP, and RPC traffic.

If you are an advanced user with a complex deployment, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones, networks, or VLANs. (Note that ASA FirePOWER modules cannot restrict preprocessing by VLAN.)

Note Tailoring preprocessing using custom network analysis policies—especially multiple network analysis policies—is an advanced task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful to allow the network analysis and intrusion policies examining a single packet to complement each other.

Benefits of custom intrusion policies

Custom intrusion policies are security configurations that

- enable tailored traffic inspection through multiple specialized intrusion policies
- manage which intrusion and inspector rules are enabled and how they are configured, and
- allow access control rules to specify which policy inspects which traffic based on multiple criteria.

In a newly created access control policy initially configured to perform intrusion prevention, the default action allows all traffic, but first inspects it with the system-provided Balanced Security and Connectivity intrusion policy. Unless you add access control rules or change the default action, all traffic is inspected by that intrusion policy.

To customize your intrusion prevention deployment, you can create multiple intrusion policies, each tailored to inspect traffic differently. Then, configure an access control policy with rules that specify which policy inspects which traffic. Access control rules can be simple or complex, matching and inspecting traffic using multiple criteria including security zone, network or geographical location, VLAN, port, application, requested URL, or user.

The main function of intrusion policies is to manage which intrusion and inspector rules are enabled and how they are configured, as follows:

- Within each intrusion policy, you should verify that all rules applicable to your environment are enabled, and improve performance by disabling rules that are not applicable to your environment. You can specify which rules should drop or modify malicious packets.
- Cisco recommendations allow you to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.

- You can modify existing rules and write new standard text rules as needed to catch new exploits or to enforce your security policies.

Other customizations you might make to an intrusion policy include:

- The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text. Note that other inspectors that detect specific threats (back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies.
- Global thresholds cause the system to generate events based on how many times traffic matching an intrusion rule originates from or is targeted to a specific address or address range within a specified time period. This helps prevent the system from being overwhelmed with a large number of events.
- Suppressing intrusion event notifications and setting thresholds for individual rules or entire intrusion policies can also prevent the system from being overwhelmed with a large number of events.
- In addition to the various views of intrusion events within the web interface, you can enable logging to syslog facilities or send event data to an SNMP trap server. For each policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events. Note that in addition to these per-policy alerting configurations, you can globally enable or disable email alerting on intrusion events for each rule or rule group. Your email alert settings are used regardless of which intrusion policy processes a packet.

Limitations of custom policies

Because preprocessing and intrusion inspection are so closely related, you **must** be careful that your configuration allows the network analysis and intrusion policies processing and examining a single packet to complement each other.

By default, the system uses one network analysis policy to preprocess all traffic handled by managed devices using a single access control policy. This diagram shows how a newly created access control policy in an inline intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.



Notice how a default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default.

You can tune preprocessing by creating and using a custom network analysis policy as the default. However, if you disable a inspector in a custom network analysis policy but the system needs to evaluate preprocessed packets against an enabled intrusion or inspector rule, the system automatically enables and uses the inspector although it remains disabled in the network analysis policy web interface.

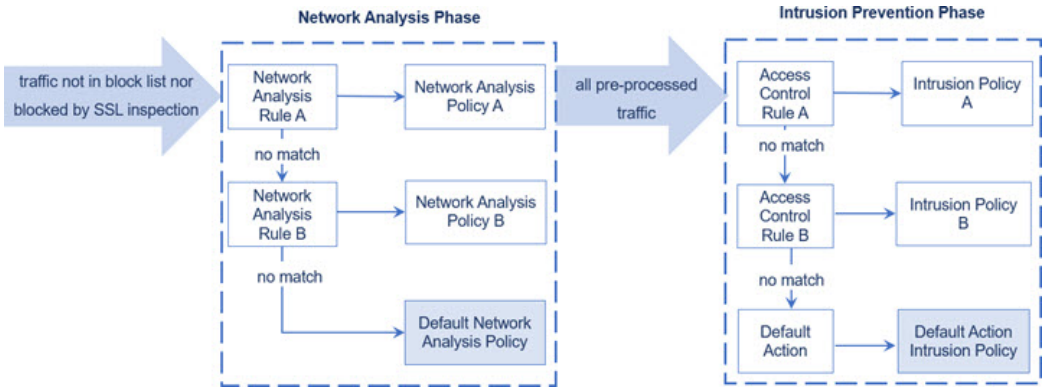
Note In order to get the performance benefits of disabling an inspector, you **must** make sure that none of your intrusion policies have enabled rules that require that inspector.

An additional challenge arises if you use multiple custom network analysis policies. For advanced users with complex deployments, you can tailor preprocessing to specific security zones, networks, and VLANs by assigning custom network analysis policies to preprocess matching traffic. (Note that ASA FirePOWER cannot restrict preprocessing by VLAN.) To accomplish this, you add custom *network analysis rules* to your access control policy. Each rule has an associated network analysis policy that governs the preprocessing of traffic that matches the rule.

Tip You configure network analysis rules as an advanced setting in an access control policy. Unlike other types of rules, network analysis rules invoke network analysis policies, rather than being contained by them.

The system matches packets to any configured network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rule is preprocessed by the default network analysis policy. While this allows you a great deal of flexibility in preprocessing traffic, keep in mind that all packets, **regardless** of which network analysis policy preprocessed them, are subsequently matched to access control rules—and thus to potential inspection by intrusion policies—in their own process. In other words, preprocessing a packet with a particular network analysis policy does **not** guarantee that the packet will be examined with any particular intrusion policy. You **must** carefully configure your access control policy so it invokes the correct network analysis and intrusion policies to evaluate a particular packet.

The diagram shows in detail how the network analysis policy (preprocessing) selection phase occurs before and separately from the intrusion prevention (rules) phase. For simplicity, the diagram eliminates the discovery and file or malware inspection phases. It also highlights the default network analysis and default-action intrusion policies.



In this scenario, an access control policy is configured with two network analysis rules and a default network analysis policy:

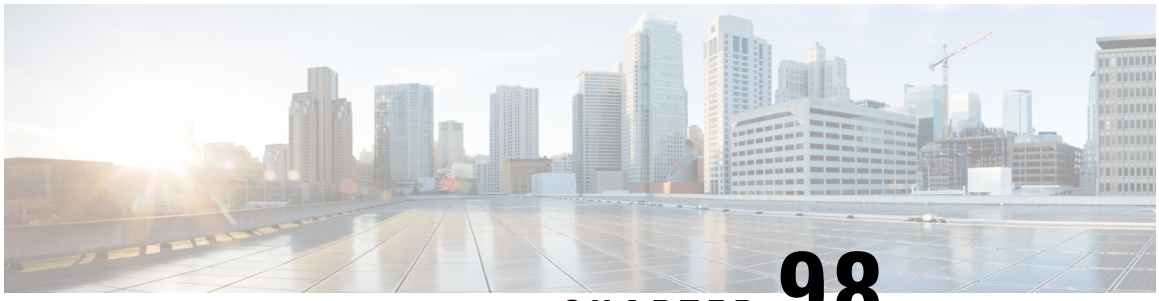
- Network Analysis Rule A preprocesses matching traffic with Network Analysis Policy A. Later, you want this traffic to be inspected by Intrusion Policy A.
- Network Analysis Rule B preprocesses matching traffic with Network Analysis Policy B. Later, you want this traffic to be inspected by Intrusion Policy B.
- All remaining traffic is preprocessed with the default network analysis policy. Later, you want this traffic to be inspected by the intrusion policy associated with the default action of the access control policy.

After the system preprocesses traffic, it can examine the traffic for intrusions. The diagram shows an access control policy with two access control rules and a default action:

- Access Control Rule A allows matching traffic. The traffic is then inspected by Intrusion Policy A.
- Access Control Rule B allows matching traffic. The traffic is then inspected by Intrusion Policy B.
- The default action of the access control policy allows matching traffic. The traffic is then inspected by the default action's intrusion policy.

Each packet's handling is governed by a network analysis policy and intrusion policy pair, but the system does **not** coordinate the pair for you. Consider a scenario where you misconfigure your access control policy so that Network Analysis Rule A and Access Control Rule A do not process the same traffic. For example, you could intend the paired policies to govern the handling of traffic on a particular security zone, but you mistakenly use different zones in the two rules' conditions. This could cause traffic to be incorrectly preprocessed. For this reason, tailoring preprocessing using network analysis rules and custom policies is an **advanced** task.

Note that for a single connection, although the system selects a network analysis policy before an access control rule, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.



CHAPTER 98

Migrate from Snort 2 to Snort 3

Starting with Version 7.0, Snort 3 is the default inspection engine for new Firewall Threat Defense deployments with Cloud-Delivered Firewall Management Center. If you are still using the Snort 2 inspection engine, switch to Snort 3 now for improved detection and performance.

Upgrading Firewall Threat Defense to Version 7.2 through 7.6 also upgrades eligible Snort 2 devices to Snort 3. For devices that are ineligible because they use custom intrusion or network analysis policies, manually upgrade to Snort 3 as described here.

Although you can switch individual devices back, you should not. Snort 2 will be deprecated in a future release and will eventually prevent Firewall Threat Defense upgrade.

- Snort 3 inspection engine, on page 2799
- Snort 2 versus Snort 3, on page 2800
- Migrating from Snort 2 to Snort 3, on page 2800
- View Snort 2 and Snort 3 base policy mapping, on page 2804
- Synchronize Snort 2 rules with Snort 3, on page 2805
- Deploy configuration changes, on page 2806
- Examples for migration, on page 2808

Snort 3 inspection engine

Snort 3 is a network inspection engine that

- serves as the default inspection engine for newly registered Firewall Threat Defense devices on version 7.0 or later,
- requires explicit enablement after upgrading Firewall Threat Defense devices to version 7.0 or later, and
- activates the Snort 3 version of the intrusion policy, which is applied through access control policies, for all traffic passing through the device.

For Firewall Threat Defense devices on version 6.x or earlier, Snort 2 is the default inspection engine.

Version switching and rule mapping

You can switch Snort versions when required. Snort 2 and Snort 3 intrusion rules are mapped and the mapping is system-provided. However, you may not find a one-to-one mapping of all the intrusion rules in Snort 2 and Snort 3. If you change the rule action for one rule in Snort 2, that change will not be retained if you switch to

Snort 3 without first synchronizing Snort 2 with Snort 3. For more information on synchronization, see Synchronize Snort 2 rules with Snort 3, on page 2805.

Snort 2 versus Snort 3

Snort 3 has been architecturally redesigned to inspect more traffic with equivalent resources when compared to Snort 2. Snort 3 provides simplified and flexible insertion of traffic parsers. Snort 3 also provides new rule syntax that makes rule writing easier and shared object rule equivalents visible.

This table lists the differences between the Snort 2 and the Snort 3 versions in terms of the inspection engine capabilities.

Feature	Snort 2	Snort 3
Packet threads	One per process	Any number per process
Configuration memory use	Number of processes * *x* GB	*x* gigabytes (GB) in total; more memory available for packets
Configuration reload	Slower	Faster; one thread can be pinned to separate cores
Rule syntax	Inconsistent and requires line escapes	Uniform system with arbitrary whitespace
Rule comments	Comments only	#, #begin and #end marks; C language style

Additional reference: Differences between Snort 2 and Snort 3 in Firepower.

Migrating from Snort 2 to Snort 3

Migrating from Snort 2 to Snort 3 requires you to switch the inspection engine of the Firewall Threat Defense device from Snort 2 to Snort 3.

Summary

Depending on your requirements, go through the stages listed in the workflow to complete the migration of your device from Snort 2 to Snort 3.

Workflow

These are the stages of migrating from Snort 2 to Snort 3:

1. Enable Snort 3 on the device.
   - Enable Snort 3 on an individual device, on page 2801
   - Enable Snort 3 on multiple devices, on page 2802
2. Convert Snort 2 custom rules to Snort 3.

- Convert all Snort 2 custom rules across all intrusion policies to Snort 3, on page 2803
- Convert Snort 2 custom rules of a single intrusion policy to Snort 3, on page 2804

3. Synchronize Snort 2 rules with Snort 3.
   - Synchronize Snort 2 rules with Snort 3, on page 2805

Prerequisites for migrating from Snort 2 to Snort 3

Consider the recommended prerequisites before migrating your device from Snort 2 to Snort 3.

- Have a working knowledge of Snort. To learn about the Snort 3 architecture, see Snort 3 Adoption.
- Back up your management center. See Backup the Management Center.
- Back up your intrusion policy. See Exporting Configurations.
- Clone your intrusion policy. To do this, you can use an existing policy as the base policy to create a copy of your intrusion policy. In the **Intrusion Policies** page, click **Create Policy** and choose an existing intrusion policy from the **Base Policy** dropdown list.

Enable Snort 3 on an individual device

Enable Snort 3 on an individual device to take advantage of enhanced security capabilities.

Use this procedure when you need to upgrade a specific device from Snort 2 to Snort 3 inspection engine. The system will automatically convert policy configurations during deployment to ensure compatibility with the selected Snort version.

Important During the deployment process, there could be a momentary traffic loss because the current inspection engine needs to be shut down.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the device to go to the device home page.

Note
The device is marked as Snort 2 or Snort 3, showing the current version on the device.

Step 3 Click the **Device** tab.

Step 4 In the Inspection Engine section, click **Upgrade**.

Note
In case you want to disable Snort 3, click **Revert to Snort 2** in the Inspection Engine section.

Step 5 Click **Yes**.

What to do next

Deploy the changes on the device. See Deploy configuration changes, on page 2806.

The system converts your policy configurations during the deployment process to make them compatible with the selected Snort version.

Enable Snort 3 on multiple devices

This task enables Snort 3 on multiple devices to provide enhanced intrusion detection capabilities.

To enable Snort 3 on multiple devices, ensure all the required Firewall Threat Defense devices are on version 7.0 or later.

Important During the deployment process, there could be a momentary traffic loss because the current inspection engine needs to be shut down.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Select all the devices on which you want to enable or disable Snort 3.

Note
The devices are marked as Snort 2 or Snort 3, showing the current version on the device.

Step 3 Click the **Select Bulk Action** dropdown list and choose **Upgrade to Snort 3**.

Step 4 Click **Yes**.

What to do next

Deploy the changes on the device. See Deploy configuration changes, on page 2806.

Convert Snort 2 Custom IPS Rules to Snort 3

Snort 2 to Snort 3 custom IPS rule conversion is a migration process that transforms custom intrusion prevention system rules from Snort 2 format to Snort 3 format.

Conversion requirements and resources

If you are using a rule set from a third-party vendor, contact that vendor to confirm that their rules successfully convert to Snort 3 or to obtain a replacement rule set written natively for Snort 3. If you have custom rules that you have written yourself, familiarize yourself with writing Snort 3 rules prior to conversion so that you

can update your rules to optimize Snort 3 detection after conversion. See the links to learn more about writing rules in Snort 3.

- https://blog.snort.org/2020/08/how-rules-are-improving-in-snort-3.html
- https://blog.snort.org/2020/10/talos-transition-to-snort-3.html

You can refer to other blogs at https://blog.snort.org/ to learn more about Snort 3 rules.

See these procedures to convert Snort 2 rules to Snort 3 rules using the system-provided tool.

- Convert all Snort 2 custom rules across all intrusion policies to Snort 3, on page 2803
- Convert Snort 2 custom rules of a single intrusion policy to Snort 3, on page 2804



Important Snort 2 network analysis policy (NAP) settings *cannot* be copied to Snort 3 automatically. NAP settings have to be manually replicated in Snort 3.

Convert all Snort 2 custom rules across all intrusion policies to Snort 3

Convert legacy Snort 2 custom rules to the updated Snort 3 format to ensure compatibility with current intrusion detection systems.

This task migrates existing Snort 2 custom rules to Snort 3 format, allowing you to maintain your custom rule configurations while upgrading to the latest intrusion detection engine. You can either import the converted rules directly or download them for review before importing.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Ensure **All Rules** is selected in the left pane.

Step 4 Click the **Tasks** drop-down list and choose:

- **Convert Snort 2 rules and import**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and import them into Firewall Management Center as Snort 3 custom rules.
- **Convert Snort 2 rules and download**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and download them into your local system.

Step 5 Click **OK**.

Note

- If you selected **Convert and import** in the previous step, then all the converted rules are saved under a newly created rule group **All Snort 2 Converted Global** under **Local Rules**.
- If you selected **Convert and download** in the previous step, then save the rules file locally. You can review the converted rules in the downloaded file and later upload them by following the steps in Add custom rules to rule groups, on page 2846.

Refer to the video Converting Snort 2 Rules to Snort 3 for additional support and information.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Convert Snort 2 custom rules of a single intrusion policy to Snort 3

Convert Snort 2 custom rules in an intrusion policy to Snort 3 format to maintain compatibility with Snort 3 environments.

When migrating to Snort 3, existing Snort 2 custom rules need to be converted to maintain their functionality. This task allows you to convert custom rules for a specific intrusion policy.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 In the **Intrusion Policies** tab, click **Show Snort 3 Sync status**.

Step 3 Click the **Sync** icon **Snort out-of-Sync** ( ) of the intrusion policy.

Note
If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in green**Snort in-Sync** ( ). It indicates that there are no custom rules to be converted.

Step 4 Read the summary and click the **Custom Rules** tab.

Step 5 Choose:

- **Import converted rules to this policy**: To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and import them into Firewall Management Center as Snort 3 custom rules.
- **Download converted rules**: To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and download them into your local system. You can review the converted rules in the downloaded file and later upload the file by clicking the upload icon.

Step 6 Click **Re-Sync**.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

View Snort 2 and Snort 3 base policy mapping

This task allows you to view the mapping between Snort 3 and Snort 2 intrusion policies.

Note Snort 2 is not supported on threat defense version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Cloud-Delivered Firewall Management Center guide that matches your Firewall Threat Defense version.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **IPS Mapping**.

Step 4 In the **IPS Policy Mapping** dialog box, click **View Mappings** to view the Snort 3 to Snort 2 intrusion policy mapping.

Step 5 Click **OK**.

Synchronize Snort 2 rules with Snort 3

This task helps ensure that Snort 2 rule override settings and custom rules are replicated on Snort 3, allowing you to start with similar coverage when transitioning between versions.

To ensure that the Snort 2 version settings and custom rules are retained and carried over to Snort 3, the Firewall Management Center provides the synchronization functionality. Synchronization helps Snort 2 rule override settings and custom rules, which you may have altered and added over the last few months or years, to be replicated on the Snort 3 version. This utility helps to synchronize Snort 2 version policy configuration with Snort 3 version to start with similar coverage.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Cloud-Delivered Firewall Management Center guide that matches your Firewall Threat Defense version.

If the Firewall Management Center is upgraded from 6.7 or earlier to 7.0 or later version, the system synchronizes the configuration. If the Firewall Management Center is a fresh 7.0 or later version, you can upgrade to a higher version, and the system will not synchronize any content during upgrade.

Before upgrading a device to Snort 3, if changes are made in Snort 2 version, you can use this utility to have the latest synchronization from Snort 2 version to Snort 3 version so that you start with a similar coverage.

Note When you move to Snort 3, manage the Snort 3 version of the policy independently. Do not use this utility as a regular operation.

Important

- Only the Snort 2 rule overrides and custom rules are copied to Snort 3 and not the other way around. You may not find a one-to-one mapping of all the intrusion rules in Snort 2 and Snort 3. Your changes to rule actions for rules that exist in both versions are synchronized when you perform the following procedure.
- Synchronization *does not* migrate the threshold and suppression settings of any custom or system-provided rules from Snort 2 to Snort 3.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **Show Snort 3 Sync status**.

Step 4 Identify the intrusion policy that is out-of-sync.

Step 5 Click the **Sync** icon **Snort out-of-Sync** ( ).

Note
If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in green**Snort in-Sync** ( ).

Step 6 Read the summary. Download a copy of the summary if required.

Step 7 Click **Re-Sync**.

Note
- The synchronized settings will be applicable on the Snort 3 intrusion engine only if it is applied on a device, and after a successful deployment.
- Snort 2 custom rules can be converted to Snort 3 with the system-provided tool. If you have any Snort 2 custom rules, click the Custom Rules tab and follow the on-screen instructions to convert the rules. For more information, see Convert Snort 2 custom rules of a single intrusion policy to Snort 3, on page 2804.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Deploy configuration changes

Deploy configuration changes to ensure that modifications made to device configurations take effect on the affected devices.

After you change configurations, deploy them to the affected devices.

Note This topic covers the basic steps involved in deploying configuration changes. We *strongly* recommend that you refer to the *Deploy Configuration Changes* topic in the latest version of the *Cisco Secure Firewall Management Center Device Configuration Guide* to understand the prerequisites and implications of deploying the changes before proceeding with the steps.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic.

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center menu bar, click **Deploy** and choose **Deployment**.

The GUI page lists the devices with out-of-date configurations having **Pending** status.

- The **Modified By** column lists the users who have modified the policies or objects. Expand the device listing to view the users who have modified the policies for each policy listing.

  Note
  Usernames are not provided for deleted policies and objects.

- The **Inspect Interruption** column indicates if traffic inspection interruption might occur in the device during deployment.

  If this column is blank for a device, it indicates that there will be no traffic inspection interruptions on that device during deployment.

- The **Last Modified Time** column specifies the last time you made configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.
- The **Status** column provides the status for each deployment.

Step 2 Identify and choose the devices on which you want to deploy configuration changes.

- Search: Search for the device name, type, domain, group, or status in the search box.
- Expand: Click **Expand Arrow** ( ) to view device-specific configuration changes to be deployed.

  If you select a checkbox next to a device, all changes made to that device and listed for the device, are included in the deployment. However, you can use **Policy selection** ( ) to select individual policies or specific configurations to deploy while withholding the remaining changes without deploying them.

  Note
  - When the status in the **Inspect Interruption** column indicates (**Yes**) that deploying will interrupt inspection, and perhaps traffic, on a Firewall Threat Defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** ( ).

- When there are changes to interface groups, security zones, or objects, the impacted devices are shown as out-of-date on the Firewall Management Center. To ensure that these changes take effect, the policies with these interface groups, security zones, or objects, also need to be deployed along with these changes. The impacted policies are shown as out-of-date on the **Preview**page on the Firewall Management Center.

Step 3 Click **Deploy**.

Step 4 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- Deploy: Continue deploying without resolving warning conditions. You cannot proceed if the system identifies errors.
- Close: Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

Configuration changes are deployed to the selected devices and take effect according to the deployment settings.

What to do next

During deployment, if there is a deployment failure, there is a possibility that the failure may impact traffic. However, it depends on certain conditions. If there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For details, see the Deploy Configuration Changes topic in the latest version of the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Examples for migration

Migrate from Snort 2 to Snort 3

Migrating from Snort 2 to Snort 3

- involves converting and adapting the Snort 2 rules to the Snort 3 rule syntax,
- optimizes the rules for improved detection and performance, and
- allows organizations to leverage the enhanced features and capabilities of Snort 3.

Organizations with Threat Defense devices managed by the Secure Firewall Management Center can opt for a hybrid deployment approach during the migration from Snort 2 to Snort 3. This approach allows for a gradual transition and minimizes potential disruptions, if any.

Benefits of migrating to snort 3

Snort 3 provides

- **Enhanced protocol support**: Snort 3 provides improved protocol support, allowing you to detect and monitor threats across a wide range of modern protocols, including encrypted traffic.
- **Streamlined rule management**: Snort 3 offers a more user-friendly rule language and rule management system, making it easier to create, modify, and manage rules effectively.
- **Improved performance**: Snort 3 has been optimized to handle higher traffic volumes more efficiently, ensuring timely threat detection and reducing the risk of performance bottlenecks.

Sample business scenario

This sample business scenario is an illustrative situation that

- demonstrates how organizations evaluate network security inspection engine migrations,
- shows the roles and motivations of security analysts and network administrators in technology decisions, and
- highlights the benefits of upgrading from legacy to modern security monitoring solutions.

Organizational migration scenario

Alice works as a security analyst in a large organization that heavily relies on the Snort inspection engine to monitor and protect their network infrastructure. The organization has been using Snort Version 2 for several years, but they have encountered some limitations and challenges.

Bob, the network administrator, is looking to migrate from Snort 2 to Snort 3 to overcome these issues and enhance his organization's network security capabilities.

This migration will also improve network security monitoring, enhance performance, and streamline rule management.

Best practices for migrating from Snort 2 to Snort 3

Follow these best practices when migrating from Snort 2 to Snort 3 to ensure data integrity and maintain security coverage.

- Back up your intrusion policy before performing the migration. See the Export Configurations task in the *Cisco Secure Firewall Management Center Administration Guide*.
- Before upgrading a device to Snort 3, if changes are made in Snort 2, use the synchronize utility to include the latest synchronization from Snort 2 to Snort 3 so that you can start with a similar coverage. See Synchronize Snort 2 rules with Snort 3, on page 2805.
- Snort 2 custom rules are not automatically converted to Snort 3 and must be manually migrated. See Convert Snort 2 Custom IPS Rules to Snort 3, on page 2802.
- Synchronization does not migrate Snort 2 rules with thresholds or suppressions. These rules must be created again in Snort 3.

Prerequisites for migration

Complete these prerequisites before starting the migration process to ensure system readiness and data protection.

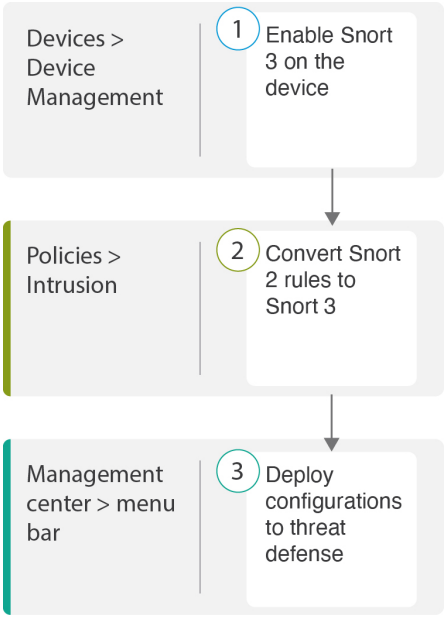
- Have a working knowledge of Snort. To learn about the Snort 3 architecture, see Snort 3 Adoption.
- Back up your management center. See Backup the Management Center.
- Back up your intrusion policy. See Exporting Configurations.

End-to-end migration workflow

Summary

Depending on your requirements, go through the stages listed in the workflow to complete the migration of your device from Snort 2 to Snort 3.

Workflow



These are the stages of migrating Snort 2 to Snort 3:

1. Enable Snort 3 on the device. See Enable Snort 3 on Threat Defense device, on page 2811.
2. Convert Snort 2 rules to Snort 3. See Convert Snort 2 rules of a single intrusion policy to Snort 3, on page 2812.
3. Deploy configuration. See Deploy configuration changes, on page 2806.

Enable Snort 3 on Threat Defense device

Enable Snort 3 on your Threat Defense device for enhanced security capabilities.

Use this procedure to upgrade from Snort 2 to Snort 3 on your Threat Defense device.

Attention During the deployment process, there could be a momentary traffic loss because the current inspection engine needs to be shut down.

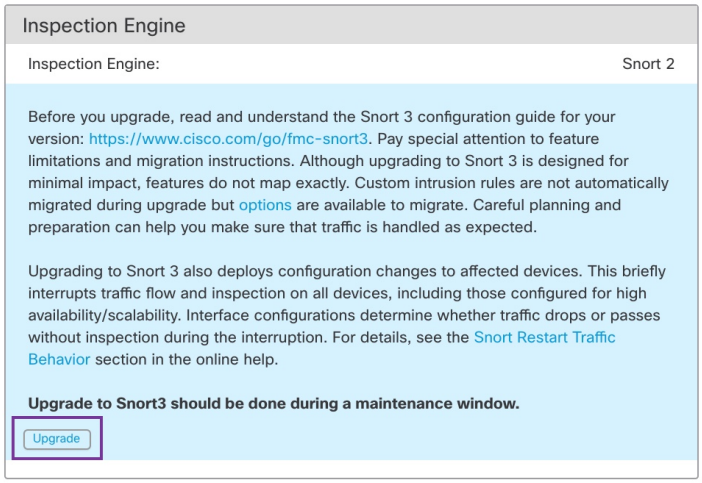
Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the corresponding device to go to the device home page.

Step 3 Click the **Device** tab.

Step 4 In the **Inspection Engine** section, click **Upgrade**.



Step 5 Click **Yes**.

What to do next

Deploy the changes on the device. See Deploy configuration changes, on page 2806.

The system converts your policy configurations during the deployment process to make them compatible with the selected Snort version.

Convert Snort 2 rules of a single intrusion policy to Snort 3

This task synchronizes and converts Snort 2 intrusion policy rules to Snort 3 format, ensuring that custom rules, thresholds, and suppressions are properly migrated to maintain security policy effectiveness when using Snort 3 engines.

When intrusion policies display an orange arrow indicator, the Snort 2 and Snort 3 versions are not synchronized. The synchronization process uses the snort2Lua tool to convert rules and may require manual intervention for custom rules, thresholds, and suppressions that cannot be automatically migrated.

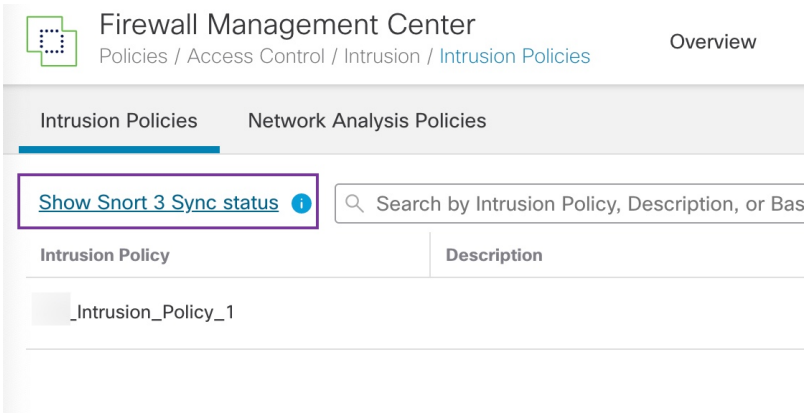
Before you begin

Follow these steps to convert Snort 2 rules of a single intrusion policy to Snort 3:

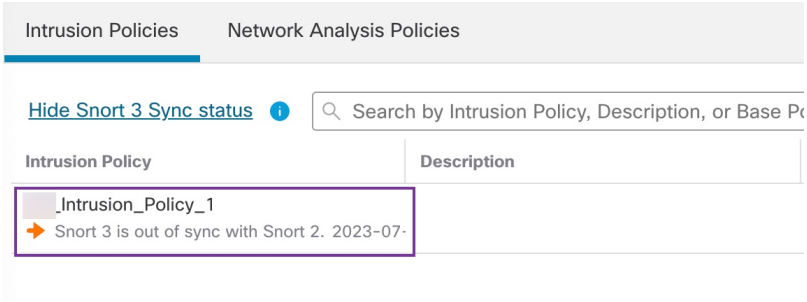
Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion** > **Intrusion Policies**.

a) In the **Intrusion Policies** tab, click **Show Snort 3 Sync status**.

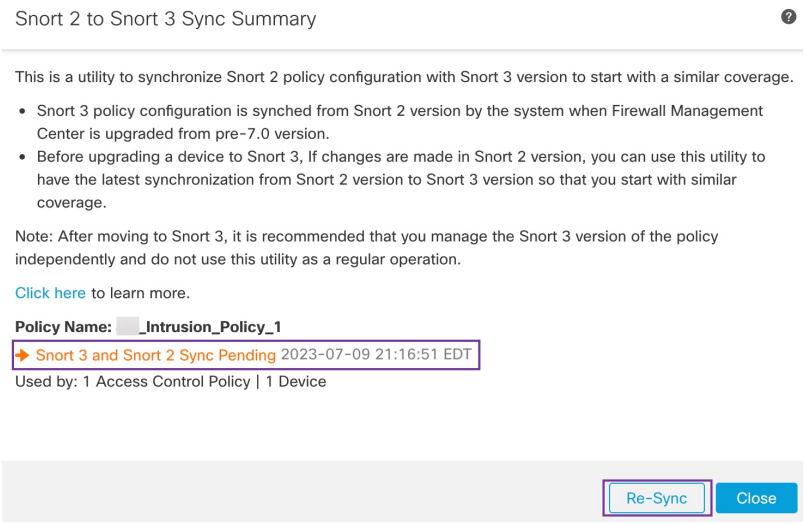


If your policy displays an orange arrow, it indicates that the Snort 2 and the Snort 3 versions of the intrusion policy are not synchronized.



b) Click the orange arrow.

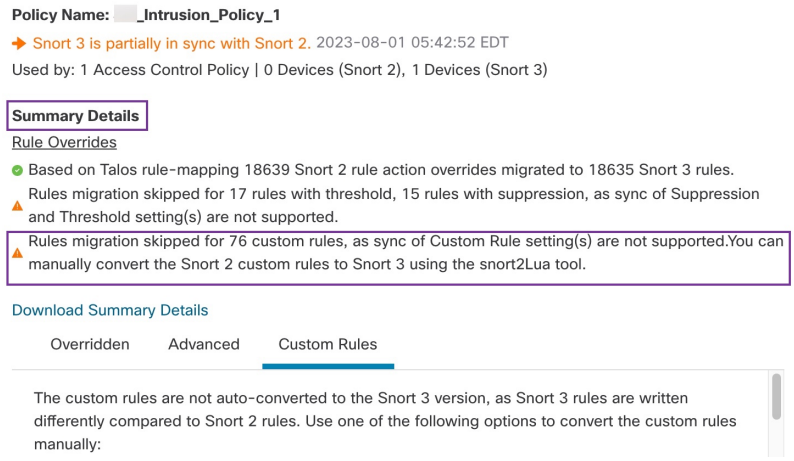
The **Snort 2 to Snort 3 Sync Summary** page displays that the Snort 2 to Snort 3 sync is pending.



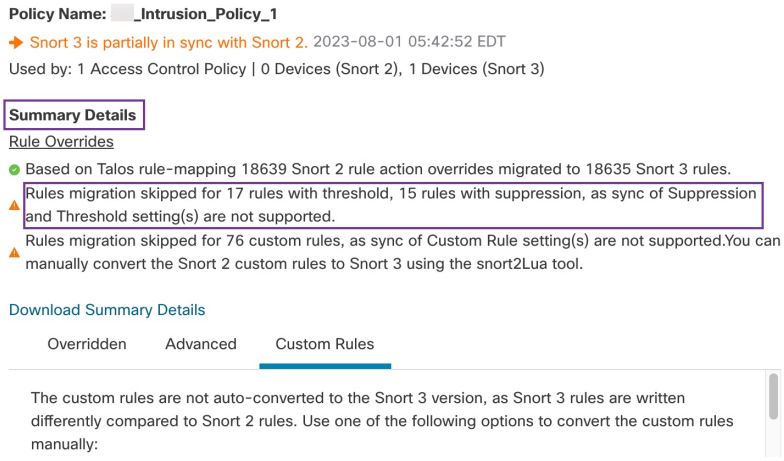
c) Click **Re-Sync** to start the synchronization.

Note
When you click **Re-Sync**, the snort2Lua tool converts the rules from Snort 2 to Snort 3.

The **Summary Details** section lists the rules that were migrated or skipped. In our use case, there are 76 custom Snort 2 rules, 17 rules with thresholds, and 15 rules with suppression that were skipped during the sync process. To migrate the custom rules, go to the next step.



To migrate rules with thresholds and suppressions, go to Step 6.



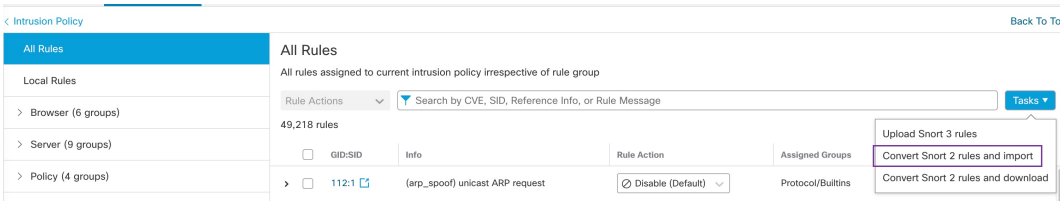
Step 2 To migrate the 76 custom rules, perform either one of these steps:

- In the **Custom Rules** tab, click the **Import** icon to convert and auto-import the local rules to the Snort 3 version of the policy.

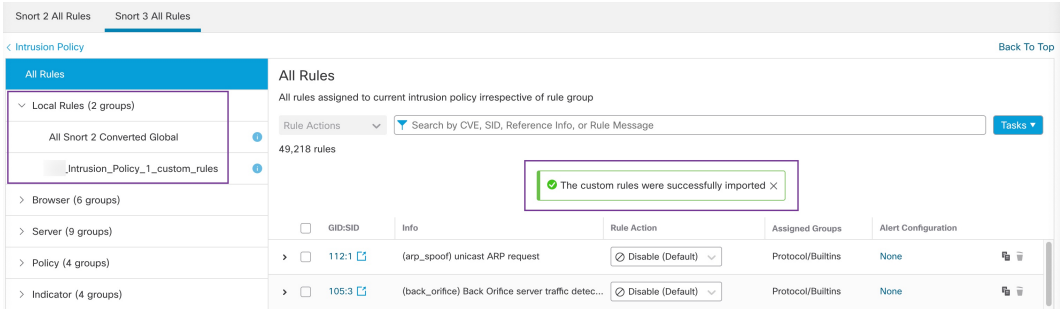


A confirmation message is displayed after the rules are successfully imported.

- Choose **Objects** > **Intrusion Rules** and click **Snort 3 All Rules**.
  - **a.** Click **Local Rules** in the left panel to check if any rules have been migrated. Notice that no custom rules from Snort 2 have been migrated.
  - **b.** From the **Tasks** drop-down list, choose **Convert Snort 2 rules and import**.

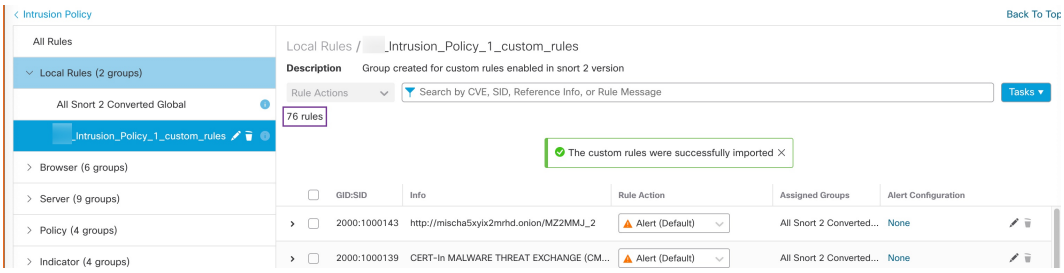


  - **c.** Click **OK**.



A rule group named (**All Snort 2 Converted Global**) is now created under **Local Rules** in the left panel.

Notice that all 76 custom rules have been migrated, as shown in the following figure.



Alternatively, you can select the **Convert Snort 2 rules and download** in the previous step to save the rules file locally. You can review the converted rules in the downloaded file and later upload them using the **Upload Snort 3 rules** option.

Step 3 Click the **Download Summary Details** link to download the rules in .txt format.

This is a sample of the summary that is displayed.

Example:


```
  "id": "00505691-15DC-0ed3-0000-004294988561",
  "name": "_Intrusion_Policy_1",
  "type": "IntrusionPolicy",
  "syncStatus": {
    "source": {
      "id": "bdce2d6a-1ebe-11ee-8e88-220032eb1fb5",
      "type": "IntrusionPolicy"
    },
    "status": "WARN",
    "description": "Migration is partially successful. Some of the rules are not copied to
 Snort3.",
    "timestamp": 1690883954814,
    "lastUser": {
      "name": "admin"
    },
    "details": [
      {
        "type": "Summary",
        "status": "INFO",
        "description": "Based on Talos rule-mapping 18639 Snort 2 rule action overrides
migrated to 18635 Snort 3 rules."
      },
      {
```

```
        "id":
"1:1000156=alert,1:1000114=alert,1:1000160=alert,1:1000135=alert,1:1000115=alert,1:1000118=alert,
1:1000092=alert,1:1000139=alert,1:1000123=alert,1:1000159=alert,1:1000149=disabled,1:1000167=alert,
1:1000133=alert,1:1000095=alert,1:1000143=alert,1:1000106=alert,1:1000153=alert,1:1000097=alert,1:1000141=alert,
1:1000148=alert,1:1000090=alert,1:1000119=alert,1:1000112=alert,1:1000138=alert,1:1000128=alert,1:1000132=alert,
1:1000134=alert,1:1000145=disabled,1:1000110=disabled,1:1000107=alert,1:1000163=alert,1:1000124=alert,1:1000125=alert,
1:1000094=alert,1:1000113=disabled,1:1000147=alert,1:1000161=alert,1:1000105=disabled,1:1000140=alert,1:1000111=alert,
1:1000102=alert,1:1000129=disabled,1:1000108=alert,1:1000144=disabled,1:1000088=alert,1:1000091=alert,1:1000131=alert,
1:1000157=alert,1:1000120=alert,1:1000126=alert,1:1000165=alert,1:1000146=alert,1:1000162=alert,1:1000116=alert,1:1000142=alert,
1:1000170=disabled,1:1000169=alert,1:1000104=alert,1:1000099=disabled,1:1000171=alert,1:1000093=alert,1:1000087=alert,1:1000100=alert,
1:1000137=alert,1:1000158=alert,1:1000103=alert,1:1000098=alert,1:1000127=disabled,1:1000130=alert,1:1000164=alert,1:1000089=alert,
         1:1000109=alert,1:1000136=alert,1:1000117=alert,1:1000166=alert,1:1000168=alert",
        "type": "PolicyInfo",
        "description": "Corresponding Snort 2 policy overridden custom (local) rules."
      },
      {
        "type": "AssignedDevices",
        "status": "INFO",
        "description": "Snort3:0 , Snort2:0"
      },
      {
        "id": "122:6",
        "type": "Threshold",
        "status": "ERROR",
        "description": "PSNG_TCP_FILTERED_DECOY_PORTSCAN"
      },
      {
        "id": "122:15",
        "type": "Threshold",
        "status": "ERROR",
        "description": "PSNG_IP_PORTSWEEP_FILTERED"
       },
      {
        "id": "122:1",
        "type": "Threshold",
        "status": "ERROR",
        "description": "PSNG_TCP_PORTSCAN"
      },
```


Step 4 Click **Close** to close the **Sync Summary** dialog box.

Step 5 To check the rules with status: ERROR, choose **Policies** > **Access Control heading** > **Intrusion** and click the **Snort 2** version of the intrusion policy.

Step 6 Under **Policy Information**, click **Rules** and filter for the rule. For example, enter **PSNG_TCP_PORTSCAN** in the **Filter**field to find the rule.

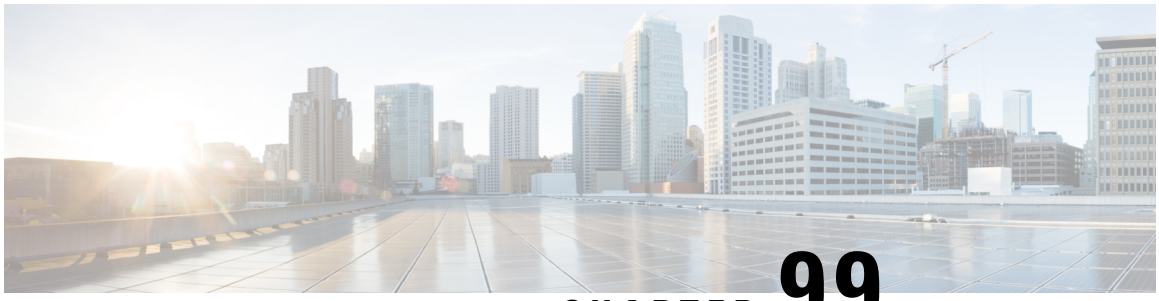
Step 7 Click **Show Details** to view the detailed version of the rule.

Step 8 Create the rule again in Snort 3 using Snort 3 rule guidelines and save the file as a .txt or .rules file. For more information, see www.snort3.org.

Step 9 Upload the custom rule that you just created locally to the list of all the Snort 3 rules. See Add Custom Rules to Rule Groups.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.



CHAPTER 99

Intrusion Policies

This chapter provides information on managing Snort 3 intrusion policies and access control rule configurations for intrusion detection and prevention.

- Intrusion policies, on page 2819
- Create a custom snort 3 intrusion policy, on page 2820
- Edit Snort 3 Intrusion Policies, on page 2821
- Change the base policy of an intrusion policy, on page 2826
- View Snort 2 and Snort 3 base policy mapping, on page 2827
- Synchronize Snort 2 rules with Snort 3, on page 2827
- Manage intrusion policies, on page 2829
- Access control rule configuration to perform intrusion prevention, on page 2830
- Tune Intrusion Policies using Rules, on page 2831
- Recommended rules, on page 2850
- Mitigate threats using MITRE framework in Snort 3 intrusion policies, on page 2854

Intrusion policies

An *intrusion policy* is a defined set of intrusion detection and prevention configurations that

- inspects traffic for security violations,
- can block or alter malicious traffic in inline deployments, and
- serves as the system's last line of defense before traffic is allowed to its destination.

At the heart of each intrusion policy are the intrusion rules. An enabled rule causes the system to generate intrusion events for (and optionally block) traffic matching the rule. Disabling a rule stops processing of the rule.

The system delivers several base intrusion policies, which enable you to take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and inspector rule states (enabled or disabled), as well as provides the initial configurations for other advanced settings.

Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

If you create a custom intrusion policy, you can:

- Tune detection by enabling and disabling rules, as well as by writing and adding your own rules.
- Use Secure Firewall recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.

An intrusion policy can drop matching packets and generate intrusion events. To configure an intrusion or preprocessor drop rule, set its state to Block.

When tailoring your intrusion policy, especially when enabling and adding rules, keep in mind that some intrusion rules require that traffic first be decoded or preprocessed in a certain way. Before an intrusion policy examines a packet, the packet is preprocessed according to configurations in a network analysis policy. If you disable a required inspector, the system automatically uses it with its current settings, although the inspector remains disabled in the network analysis policy web interface.

Caution Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

After you configure a custom intrusion policy, you can use it as part of your access control configuration by associating the intrusion policy with one or more access control rules or an access control policy's default action. This forces the system to use the intrusion policy to examine certain allowed traffic before the traffic passes to its final destination. A variable set that you pair with the intrusion policy allows you to accurately reflect your home and external networks and, as appropriate, the servers on your network.

By default, the system disables intrusion inspection of encrypted payloads. This default reduces false positives and improves performance when an encrypted connection matches an access control rule with intrusion inspection enabled.

Refer to the video for additional support and information - Snort 3 Intrusion Policy Overview.

Create a custom snort 3 intrusion policy

Create a custom Snort 3 intrusion policy to tailor intrusion detection and prevention settings to your network security requirements.

Snort 3 intrusion policies define how the system inspects traffic and responds to potential threats. You can create custom policies based on existing system-provided policies or other custom policies to meet specific security needs.

Before you begin

Follow these steps to create a custom Snort 3 intrusion policy:

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Click **Create Policy**.

Step 3 Enter a unique **Name** and, optionally, a **Description**.

Step 4 Choose the **Inspection Mode**.

The selected action determines whether intrusion rules block and alert (**Prevention** mode) or only alert (**Detection** mode).

Note
Before selecting the prevention mode, you might want block rules to alert only so you can identify rules that cause many false positives.

Step 5 Choose the **Base Policy**.

You can use either a system-provided policy or an existing policy as your base policy.

Step 6 Click **Save**.

The new policy has the same settings as its base policy.

What to do next

To customize the policy, see Edit Snort 3 Intrusion Policies, on page 2821.

Edit Snort 3 Intrusion Policies

Use this task to understand how to access, modify, and manage Snort 3 intrusion policies in the Firewall Management Center, including editing policy settings, managing rule groups, and applying recommendations to protect your network.

While editing a Snort 3 policy, all the changes are saved immediately. No additional action is required to save the changes.

Before you begin

-

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion** .

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **Snort 3 Version** next to the intrusion policy you want to configure.

Step 4 Edit your policy:

- Change the mode—Click the **Mode** drop-down to change the inspection mode.

Caution
The inspection mode is changed only for the Snort 3 version of the policy. The existing inspection mode is retained in the Snort 2 version as is, which means that your Snort 2 and Snort 3 versions of the policy will have different inspection modes. We recommend that you use this option with caution.

- **Prevention** —Triggered Block rules create an event (alert) and drop the connection.
- **Detection** —Triggered Block rules create an alert.

  You can choose the detection mode before going for prevention. For example, before choosing the prevention mode, you might want block rules to alert only, so that you can identify rules that cause a lot of false positives.

Step 5 Click the **Base Policy** layer that defines the intrusion policy's default settings.

- Search rules—Use the search field to filter the display. You can enter the GID, SID, rule message, or reference info. For example, GID:1; SID:9621—to display only rule 1:962, SID:9621,9622,9623—to display multiple rules with different SIDs. You can also click inside the Search text box to choose any of the following options:
  - apply the filters **Action = Alert** , or **Action: Block**
  - apply the **Disabled Rules** filter
  - show **Custom/User Defined Rules**
  - filter by GID, SID, or GID:SID
  - filter by CVE
  - filter by comment
- View filtered rules—Click any of the **Presets** to view rules that are set to alert, block, disabled, and so on.

  Overridden rules indicate the rules where the rule action has been changed from the default action to a different action. Note that, once changed, the rule action status is Overridden even if you change it back to its original default action. However, if you select **Revert to default** from the **Rule Action** drop-down list, the Overridden status is removed.

  Advanced Filters provides filter options based on the Lightweight Security Package (LSP) releases, Classifications of intrusions, and Microsoft Vulnerabilities.
- View rule documentation—Click the rule ID or the **Rule Documentation** icon to display Talos documentation for the rule.
- View a rule details—Click the **Expand Arrow** ( ) icon in a rule row to view the rule details.
- Add rule comments—Click **Comment** ( ) under the Comments column to add comments for a rule.

Step 6 **Group Overrides** —Click the **Group Overrides** layer that lists all the categories of rule groups. The top-level parent rule group with Description, Overrides, Enabled Groups, and other details is displayed. Parent rule groups cannot be updated and are read-only. Only the leaf rule groups can be updated. In each rule group,

you can traverse up to the last leaf group. Across each group, you can override, include, and exclude rule groups. In the leaf rule groups, you can:

- Search rule groups—Use the search field to enter keywords and search for rule groups.
- In the left panel, you can choose any of the preset filter options to search for rule groups:
  - All—For displaying all rule groups.
  - Excluded—For excluded groups.
  - Included—For included groups.
  - Overridden—For rule group configuration that is overridden.
- Set the security level for a rule group—Navigate to the required rule group on the left pane and click it. Click **Edit** next to the **Security Level** of the rule group to increase or decrease the security level based on system-defined rule settings.

  In the **Edit Security Level** dialog box, you have the option to click **Revert to Default** , which reverts the changes you made.

  The Firewall Management Center automatically changes the action for the rules of the rule group for the configured security level. In the **Rule Overrides** layer, notice the count of Block Rules and Disabled Rules in the **Presets** every time you change the security level.
- You can make bulk changes to the security level to change the security level of all rule groups within a particular rule category. Bulk security level applies to rule groups that have more than one rule group. After a bulk update of rule groups, you can still update the security level of any of the associated rule groups within it.

  There can be **mixed** security levels within rule groups; **mixed** indicates that the child groups contain a mix of security levels within the parent rule group.
- Include or exclude rule groups—The rule groups displayed are the default rule groups associated with the system-provided base intrusion policy. You can include and exclude rule groups from the intrusion policy. An excluded rule group is removed from the intrusion policy and its rules are not applied on the traffic. For information on uploading custom rules in Firewall Management Center , see Add custom rules to rule groups, on page 2846 .

  To exclude a rule group:

  a. Navigate the Rule Groups pane and choose the rule group that you want to exclude.

  b. Click the **Exclude** hyperlink on the right-pane.

  c. Click **Exclude** .

  To include a new rule group or multiple rule groups with the uploaded custom rules or a previously excluded rule group:

  a. Click **Add** (+) next to the rule group filter dropdown list.

  b. Choose all the rule groups you want to add by checking the check box next to it.

  c. Click **Save** .

- For a leaf rule group, click the icon under the **Override** column header to see the rule action trail, which describes the sequence of overridden rule actions that can be assigned due to the base policy and group overrides for an intrusion rule. Rule actions can be obtained from either the base policy configurations or the user group override. The user group override takes the priority between the two; priority refers to the final overridden action that is assigned to the rule group.
- Click the rule count (number) under the **Rule Count** column header to see a summary of rules that are part of the rule group.

Step 7 **Recommendations** —Click the **Recommendations** layer if you want to generate and apply Cisco recommended rules. Recommendations use the host database to enable or disable rules, based on known vulnerabilities.

Step 8 **Rule Overrides** —Click the **Rule Overrides** layer to choose any of the presets to view rules, which are set to alert, block, disabled, overridden, rewrite, pass, drop, or reject.

- The **Set By** column shows the default set by state (Base Policy) or modified rule state by Group Overrides, Rule Overrides, or Recommendations. The **Set By** column in **All Rules** (in the left pane) shows the trail of rule action override actions based on priority order. The priority order of rule actions is Rule Override > Recommendations > Group Override > Base Policy.
- Modify **Rule Action** —To modify rule actions, choose either of the following:
  - Bulk edit—Choose one or more rules, then choose the required action from the **Rule Action** drop-down list; and click **Save** .

    Note
    Bulk rule action changes are supported only for the first 500 rules.

  - Single rule edit—Choose the action for the rule from the drop-down list in the **Rule Action** column.

  Rule actions are:

  - **Block** — Generates event, blocks current matching packet and all the subsequent packets in this connection.
  - **Alert** — Generates only events for matching packet and does not drop packet or connection.
  - **Disable** —Does not match traffic against this rule. No events are generated.
  - **Revert to default** —Reverts to the system default action.
  - **Pass** — No events are generated, allows packet to pass without further evaluation by any subsequent Snort rules.

    Note
    The Pass action is available only for custom rules and not for system-provided rules.

  - **Drop** — Generates event, drops matching packet and does not block further traffic in this connection.
  - **Reject** — Generates event, drops matching packet, blocks further traffic in this connection and sends TCP reset if it is a TCP protocol to source and destination hosts.

    Behavior of reject in different firewall modes and IP address or source or destination in relation to Client or Server : Snort sends RST packets to both client and server in cases of routed, inline, and bridged interfaces. Snort sends two RST packets. The RST packet in the client's direction will have the source set to the server's IP and the destination set to client's IP. The RST packet in the server's direction will have the source set to the client's IP and the destination set to the server's IP.

- **Rewrite** — Generates event and overwrites packet contents based on the replace option in the rule.

For IPS rule action logging, see Rule action logging, on page 2825 .

If there is a **React** rule, it is converted to an alert action.

Step 9 Click the **Summary** layer for a holistic view of the current changes to the policy. The policy summary page contains the following information:

- Rule distribution of the policy, that is, active rules, disabled rules, and similar categories.
- Option to export policy and generate report of the intrusion policy.
- Base policy details.
- Option to generate recommendations.
- Group overrides list the groups you have overridden.
- Rule overrides that shows the list of rules that you have overridden.
- In the **Summary** layer, click the ? icon to open a pop-up window that explains the Snort layering concepts.

To change the base policy, see Change the base policy of an intrusion policy, on page 2826 .

Note
You can navigate to **Objects** > **Intrusion Rules** and click the **Snort 3 All Rules** tab and traverse through all the intrusion rule groups. The parent rule group lists the associated child groups and rule count.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806 .

Rule group reporting

Rule group reporting feature reflects rule groups in generated intrusion events, includes MITRE tactics and techniques identification, and provides dedicated columns for both MITRE and non-MITRE rule groups in intrusion event displays.

To view the intrusion events in Cloud-Delivered Firewall Management Center, choose **Analysis** > **Unified Events**. In the **Unified Events** page, explicitly enable the **MITRE ATT&CK** column for rule group reporting.

Rule action logging

From Cloud-Delivered Firewall Management Center 7.2.0 onwards, in the **Intrusion Events** window, the event in the **Inline Result** column displays the same name as the IPS action applied to the rule, so that you can see the action that was applied on the traffic matching the rule.

For the IPS actions, the tables show the events that are displayed in the **Inline Result** column of the **Intrusion Events** window and **Action** column for **Intrusion Event Type** in the **Unified Events** window.

IPS Action for Snort 3	Inline Result - Cloud-Delivered Firewall Management Center 7.1.0 and earlier	Inline Result -Cloud-Delivered Firewall Management Center 7.2.0 onwards
Alert	Pass	Alert
Block	Dropped/Would Have Dropped/Partially Dropped	Block/Would Block/Partial Block
Drop	Dropped/Would have dropped	Drop/Would drop
Reject	Dropped/Would have dropped	Reject/Would reject
Rewrite	Allow	Rewrite

Important

- In case of a rule without the "Replace" option, the **Rewrite** action is displayed as **Would Rewrite**.
- The **Rewrite** action would also be displayed as **Would Rewrite** if the "Replace" option is specified, but the IPS policy is in Detection mode or the device is in Inline-TAP or Passive mode.

Note In case of backward compatibility (Cloud-Delivered Firewall Management Center 7.2.0 managing a Firewall Threat Defense 7.1.0 device), the events mentioned are applicable only to the Alert IPS action where **Pass** is displayed as **Alert** for events. For all the other actions, the events for Cloud-Delivered Firewall Management Center 7.1.0 are applicable.

Change the base policy of an intrusion policy

Change the base policy to use a different system-provided or custom policy as the foundation for your intrusion policy configuration.

You can choose a different system-provided or custom policy as your base policy. You can chain up to five custom policies, with four of the five using one of the other four previously created policies as its base policy; the fifth must use a system-provided policy as its base.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Click **Edit** ( ) next to the intrusion policy you want to configure.

Step 3 Choose a policy from the **Base Policy** drop-down list.

Step 4 Click **Save**.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

View Snort 2 and Snort 3 base policy mapping

This task allows you to view the mapping between Snort 3 and Snort 2 intrusion policies.

Note Snort 2 is not supported on threat defense version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Cloud-Delivered Firewall Management Center guide that matches your Firewall Threat Defense version.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **IPS Mapping**.

Step 4 In the **IPS Policy Mapping** dialog box, click **View Mappings** to view the Snort 3 to Snort 2 intrusion policy mapping.

Step 5 Click **OK**.

Synchronize Snort 2 rules with Snort 3

This task helps ensure that Snort 2 rule override settings and custom rules are replicated on Snort 3, allowing you to start with similar coverage when transitioning between versions.

To ensure that the Snort 2 version settings and custom rules are retained and carried over to Snort 3, the Firewall Management Center provides the synchronization functionality. Synchronization helps Snort 2 rule override settings and custom rules, which you may have altered and added over the last few months or years, to be replicated on the Snort 3 version. This utility helps to synchronize Snort 2 version policy configuration with Snort 3 version to start with similar coverage.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Cloud-Delivered Firewall Management Center guide that matches your Firewall Threat Defense version.

If the Firewall Management Center is upgraded from 6.7 or earlier to 7.0 or later version, the system synchronizes the configuration. If the Firewall Management Center is a fresh 7.0 or later version, you can upgrade to a higher version, and the system will not synchronize any content during upgrade.

Before upgrading a device to Snort 3, if changes are made in Snort 2 version, you can use this utility to have the latest synchronization from Snort 2 version to Snort 3 version so that you start with a similar coverage.

Note When you move to Snort 3, manage the Snort 3 version of the policy independently. Do not use this utility as a regular operation.

Important

- Only the Snort 2 rule overrides and custom rules are copied to Snort 3 and not the other way around. You may not find a one-to-one mapping of all the intrusion rules in Snort 2 and Snort 3. Your changes to rule actions for rules that exist in both versions are synchronized when you perform the following procedure.
- Synchronization *does not* migrate the threshold and suppression settings of any custom or system-provided rules from Snort 2 to Snort 3.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **Show Snort 3 Sync status**.

Step 4 Identify the intrusion policy that is out-of-sync.

Step 5 Click the **Sync** icon **Snort out-of-Sync** ( ).

Note
If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in green**Snort in-Sync** ( ).

Step 6 Read the summary. Download a copy of the summary if required.

Step 7 Click **Re-Sync**.

Note

- The synchronized settings will be applicable on the Snort 3 intrusion engine only if it is applied on a device, and after a successful deployment.
- Snort 2 custom rules can be converted to Snort 3 with the system-provided tool. If you have any Snort 2 custom rules, click the Custom Rules tab and follow the on-screen instructions to convert the rules. For more information, see Convert Snort 2 custom rules of a single intrusion policy to Snort 3, on page 2804.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Manage intrusion policies

This task enables you to perform comprehensive management operations on your intrusion policies including creation, modification, deletion, export, deployment, and reporting.

On the Intrusion Policy window (**Policies** > **Access Control heading** > **Intrusion**) you can view your current custom intrusion policies, along with the following information:

- The number of access control policies and devices using the intrusion policy to inspect traffic.
- Indicates the domain where the policy was created in a multidomain deployment.

Note Snort 2 is not supported on threat defense version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Cloud-Delivered Firewall Management Center guide that matches your Firewall Threat Defense version.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Manage your intrusion policy:

- Create — Click **Create Policy**; see Create a custom snort 3 intrusion policy, on page 2820.
- Delete — Click **Delete** ( ) next to the policy you want to delete. The system prompts you to confirm and informs you if another user has unsaved changes in the policy. Click **OK** to confirm.

  If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Edit intrusion policy details — Click **Edit** ( ) next to the policy you want to edit. You can edit the **Name**, **Inspection Mode**, and the **Base Policy** of the intrusion policy.
- Edit intrusion policy settings — Click **Snort 3 Version**; see Edit Snort 3 Intrusion Policies, on page 2821.
- Export — If you want to export an intrusion policy to import on another Firewall Management Center, click Export; see the *Exporting Configurations* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.
- Deploy — Choose **Deploy** > **Deployment**; see Deploy configuration changes, on page 2806.

- Report — Click **Report**; see the *Generating Current Policy Reports* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*. Generates two reports, one for each policy version.

Access control rule configuration to perform intrusion prevention

An access control rule configuration to perform intrusion prevention is a security mechanism that

- enables multiple access control rules within an access control policy to be associated with intrusion policies,
- permits intrusion inspection for Allow or Interactive Block access control rules to match different intrusion inspection profiles against different types of traffic, and
- evaluates traffic before it reaches its final destination on the network.

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

System-provided and custom intrusion policies

Cisco delivers several intrusion policies with the system. By using system-provided intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and preprocessor rule states, as well as provides the initial configurations for advanced settings. You can use system-provided policies as-is, or you can use them as the base for custom policies. Building custom policies can improve the performance of the system in your environment and provide a focused view of the malicious traffic and policy violations occurring on your network.

Connection and intrusion event logging

When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, it saves that event to the Management Center. The system also automatically logs the end of the connection where the intrusion occurred to the Management Center database, regardless of the logging configuration of the access control rule.

Access control rule configuration and intrusion policies

The number of unique intrusion policies you can use in a single access control policy depends on the model of the target devices; more powerful devices can handle more. Every unique **pair** of intrusion policy and variable set counts as one policy. Although you can associate a different intrusion policy-variable set pair with each Allow and Interactive Block rule (as well as with the default action), you cannot deploy an access control policy if the target devices have insufficient resources to perform inspection as configured.

Configure an access control rule to perform intrusion prevention

Configure an access control rule to enable intrusion prevention inspection on network traffic, allowing you to detect and prevent malicious activities based on intrusion signatures and policies.

You must be an Admin, Access Admin, or Network Admin to perform this task.

Follow these steps to configure an access control rule to perform intrusion prevention:

Procedure

Step 1 In the access control policy editor, create a new rule or edit an existing rule; see the *Access Control Rule Components* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.

Step 2 Ensure the rule action is set to **Allow**, **Interactive Block**, or **Interactive Block with reset**.

Step 3 Click **Inspection**.

Step 4 Choose a system-provided or a custom intrusion policy, or choose **None** to disable intrusion inspection for traffic that matches the access control rule.

Step 5 If you want to change the variable set associated with the intrusion policy, choose a value from the **Variable Set** drop-down list.

Step 6 Click **Save** to save the rule.

Step 7 Click **Save** to save the policy.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Tune Intrusion Policies using Rules

This section provides information on custom rules in Snort 3, intrusion rule action, intrusion event notification filters in an intrusion policy, converting Snort 2 custom rules to Snort 3, and adding rule groups with custom rules to an intrusion policy.

Tuning intrusion rules

Tuning intrusion rules is a configuration process that

- enables rule state modification for shared object rules, standard text rules, and inspector rules,
- controls event generation and traffic processing based on rule configuration, and
- allows filtering of rules to display specific subsets for targeted configuration changes.

Rule state configuration

You enable a rule by setting its rule state to Alert or to Block. Enabling a rule causes the system to generate events on traffic matching the rule. Disabling a rule stops processing of the rule. You can also set your intrusion policy so that a rule set to Block generates events on, and drops, matching traffic.

You can filter rules to display a subset of rules, enabling you to select the exact set of rules where you want to change rule states or rule settings.

When an intrusion rule or rule argument requires a disabled inspector, the system automatically uses it with its current configuration even though it remains disabled in the web user interface (UI) of the network analysis policy.

Note We recommend that you do not modify shared object rules and you only enable or disable these rules for your threat defense device. To create custom Snort rules, contact Cisco support.

Intrusion rule types

An intrusion rule is a specified set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities in your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule, and triggers the rule if the data packet meets all the conditions specified in the rule.

- *intrusion rules*, which are subdivided into *shared object rules* and *standard text rules*
- *inspector rules*, which are associated with a detection option of the packet decoder or with one of the inspectors included with the system.

Rule type attributes and characteristics

This table summarizes attributes of these rule types:

Table 151: Intrusion rule types

Type	Generator ID (GID)	Snort ID (SID)	Source	Can Copy?	Can Edit?
shared object rule	3	lower than 1,000,000	Cisco Talos Intelligence Group (Talos)	yes	limited
standard text rule	1 (Global domain or legacy GID)	lower than 1,000,000	Talos	yes	limited
	1,000-2,000 (descendant domain)	1,000,000 or higher	Created or imported by user	yes	yes

Type	Generator ID (GID)	Snort ID (SID)	Source	Can Copy?	Can Edit?
preprocessor rule	decoder- or preprocessor-specific	lower than 1,000,000	Talos	no	no
		1,000,000 or higher	Generated by the system during option configuration	no	no

You cannot save changes to any rule created by Talos, but you can save a copy of a modified rule as a custom rule. You can modify either variables used in the rule or rule header information (such as source and destination ports and IP addresses). In a multidomain deployment, rules created by Talos belong to the Global domain. Administrators in descendant domains can save local copies of the rules, which they can then edit.

For the rules it creates, Talos assigns default rule states in each default intrusion policy. Most preprocessor rules are disabled by default and must be enabled if you want the system to generate events for preprocessor rules and, in an inline deployment, drop offending packets.

Custom rules in Snort 3

A custom rule in Snort 3 is an intrusion detection rule that

- enables users to create custom threat detection by importing local rule files with `.txt` or `.rules` extensions,
- belongs to rule groups and can be included in multiple groups simultaneously, and
- receives a unique rule number in `GID:SID:Rev` format with system-generated identifiers.

Rule number format

When you create a custom intrusion rule, the system assigns it a unique rule number, which has the format `GID:SID:Rev`. The elements of this number are:

- **GID**—Generator ID. For custom rules, it is not necessary to specify the GID. The system automatically generates the GID based on whether you are in the Global domain or a sub-domain while uploading the rules. For all standard text rules, this value is 2000 for a Global domain.
- **SID**—Snort ID. Indicates whether the rule is a local rule of a system rule. When you create a new rule, assign a unique SID to the rule.

  SID numbers for local rules start at 1,000,000, and the SID for each new local rule is incremented by one.
- **Rev**—The revision number. For a new rule, the revision number is one. Each time you modify a custom rule the revision number should be incremented by one.

In a custom standard text rule, you set the rule header settings and the rule keywords and arguments. You can use the rule header settings to focus the rule to only match traffic using a specific protocol and traveling to or from specific IP addresses or ports.

To check whether a SID is enabled or disabled, verify the entries in the snort.lua file located in the `./file-contents/ngfw/var/sf/detection_engines/<ID>/IPS/<ID>` directory.

- If the SID is disabled by default, no entry will be present in the file.

- If the SID is manually enabled, you will see an entry with **enable:yes**.
- If the SID is disabled after being manually enabled, the entry remains in the file and will display **enable:no**.

Note

- Snort 3 custom rules cannot be edited. Ensure that custom rules have a valid classification message for `classtype` within the rule text. If you import a rule without a classification or with an incorrect classification, delete and recreate the rule.
- You can create custom intrusion rules using Snort 3. However, support for tuning and troubleshooting these rules is currently not available.
- The *classtype* in a Snort rule assigns a classification to the rule indicating the type of attack that is associated with an event. A priority level of 1-4 is also associated with each *classtype*. However, the priority level for certain *classtypes* on the Threat Defense device do not match the open-source Snort *classtype* priority levels that are mentioned in the Snort documentation. For example, *tcp-connection* has a priority of 4 in open-source Snort while a priority of 3 is assigned to it on the Threat Defense device.

Sensitive data detection

Sensitive data such as social security numbers, credit card numbers, emails, and so on may be leaked onto the internet, intentionally or accidentally. Sensitive data detection is used to detect and generate events on possible sensitive data leakage. Events are generated only when a significant amount of Personally Identifiable Information (PII) is transferred. Sensitive data detection can mask PII in the output of events.

Use the `sd_pattern` IPS option to detect and filter PII. This information includes credit card numbers, U.S. Social Security numbers, phone numbers, and email addresses. A regular expression (regex) syntax is available for defining your own PII.

The sd_pattern option has the following settings:

- Pattern—An implicit, required setting that specifies the regular expression to look for in the PDU. The regex must be written in PCRE syntax.
- Threshold—An explicit, optional setting that specifies the number of matches in the PDU required to generate an event.

  The `sd_pattern` as IPS rule option is available in Snort with no requirements for additional inspectors. The syntax for the rule option is:

  ```
  sd_pattern: "<pattern>"[, threshold <count>];
  ```

  For example:

  ```
  sd_pattern:"credit_card", threshold 2;
  ```

There are five built-in patterns for sensitive data. To use the built-in patterns in the "pattern" setting, you must specify the name of the PII type that needs to be matched and the necessary regex is substituted for it. The PII name and regex mappings or patterns are described here:

Table 152: Built-in patterns

PII Name	Pattern
credit_card	`\d{4}\D?\d{4}\D?\d{2}\D?\d{2}\D?\d{3,4}`

PII Name	Pattern
us_social	`[0-8]\d{2}-\d{2}-\d{4}`
us_social_nodashes	`[0-8]\d{8}`
Email	`[a-zA-Z0-9!#$%&'*+\/=?^_`{\|}~-]+(?:\.[a-zA-Z0-9!#$%&'*+\/=?^_`{|}~-]+)*@(?:[a-zA-Z0-9](?:[a-zA-Z0-9-]*[a-zA-Z0-9])?\.)+[a-zA-Z0-9](?:[a-zA-Z0-9-]*[a-zA-Z0-9])?`
us_phone	`(?:\+?1[-\.\s]?)?\(?([2-9][0-8]\d)\)?[-\.\s]([2-9]\d{2})[-\.\s](\d{4})`

Masking for data matching these patterns only work with system-provided rules or built-in patterns for Credit Cards, U.S. Social Security numbers, emails, and U.S. phone numbers. Masking does not work for custom rules or user-defined PII patterns. Rules are available in the Lightweight Security Package (LSP) for sensitive data, gid:13. By default, they are not enabled in any system-provided policy.

The sensitive data rules in LSP cover all built-in patterns and have the following threshold values:

- credit_card: 2
- us_social: 2
- us_social_nodashes: 20
- email: 20
- us_phone: 20

You can use the sd_pattern option to create custom rules and modify existing rules. To do this, use the Snort 3 intrusion policy interface.

An example of a rule with sd_pattern with a custom pattern and threshold:

alert tcp (sid: 1000000001; sd_pattern:"[\w-\.]+@([\w-]+\.)+[\w-]{2,4}",threshold 4; msg: "email, threshold 4")

Examples

An example of custom rules using sensitive data detection:

Rule with built-in pattern

```
alert tcp (
        msg:"SENSITIVE-DATA Email";
        flow:only_stream;
        pkt_data;
        sd_pattern:"email", threshold 5;
        service:http, smtp, ftp-data, imap, pop3;
        gid:2000;
        sid:1000001;
)
```

Rule with custom pattern

```
alert tcp (
        msg:"SENSITIVE-DATA US phone numbers";
        flow:only_stream;
        file_data;
        sd_pattern:"+?3?8?(0[\s\.-]\d{2}[\s\.-]\d{3}[\s\.-]\d{2}[\s\.-]\d{2})", threshold
2;
        service:http, smtp, ftp-data, imap, pop3;
```

```
        gid:2000;
        sid:1000002;
)
```

Here are some more examples of complete Snort IPS rules with built-in sensitive data patterns:

- alert tcp (sid:1; msg:"Credit Card"; sd_pattern:"credit_card", threshold 2;)
- alert tcp (sid:2; msg:"US Social Number"; sd_pattern:"us_social", threshold 2;)
- alert tcp (sid:3; msg:"US Social Number No Dashes"; sd_pattern:"us_social_nodashes", threshold 2;)
- alert tcp (sid:4; msg:"US Phone Number"; sd_pattern:"us_phone", threshold 2;)
- alert tcp (sid:5; msg:"Email"; sd_pattern:"email", threshold 2;)

Disabling data masking is not supported in the Secure Firewall Management Center, and Secure Firewall Device Manager.

For information on how to add custom rules to rule groups, see Add custom rules to rule groups, on page 2846.

View Snort 3 intrusion rules in an intrusion policy

This task allows you to view and interact with Snort 3 intrusion rules within an intrusion policy, enabling you to manage rule display settings and access detailed rule information.

You can adjust how rules are displayed in the intrusion policy. You can also display the details for a specific rule to see rule settings, rule documentation, and other rule specifics.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Click **Snort 3 Version** next to the policy.

Step 3 While viewing the rules, you can:

- Filter the rules.
- Choose a rule group to see rules related to that group.
- View an intrusion rule's details.
- View rule comments.
- View rule documentation.

See Edit Snort 3 Intrusion Policies, on page 2821 for details on performing these tasks.

Intrusion rule actions

An intrusion rule action is a policy configuration mechanism that

- enables or disables rules within individual intrusion policies,
- specifies which action the system takes if monitored conditions trigger the rule, and
- inherits default actions from the base policy used during intrusion policy creation.

Default action behavior

The Cisco Talos Intelligence Group (Talos) sets the default action of each intrusion and inspector rule in each default policy. For example, a rule may be enabled in the Security over Connectivity default policy and disabled in the Connectivity over Security default policy. Talos sometimes uses a rule update to change the default action of one or more rules in a default policy. If you allow rule updates to update your base policy, you also allow the rule update to change the default action of a rule in your policy when the default action changes in the default policy you used to create your policy (or in the default policy it is based on). Note, however, that if you have changed the rule action, the rule update does not override your change.

When you create an intrusion rule, it inherits the default actions of the rules in the default policy you use to create your policy.

Intrusion rule action options

In an intrusion policy, you can set a rule's action to specific values that determine how the system responds to matching traffic.

Alert

You want the system to detect a specific intrusion attempt and generate an intrusion event when it finds matching traffic. When a malicious packet crosses your network and triggers the rule, the packet is sent to its destination and the system generates an intrusion event. The malicious packet reaches its target, but you are notified through the event logging.

Block

You want the system to detect a specific intrusion attempt, drop the packet containing the attack, and generate an intrusion event when it finds matching traffic. The malicious packet never reaches its target, and you are notified through the event logging.

Disable

You do not want the system to evaluate matching traffic.

Note Choosing either the **Alert** or **Block** options enables the rule. Choosing **Disable** disables the rule.

We strongly recommend that you **do not** enable all the intrusion rules in an intrusion policy. The performance of your managed device is likely to degrade if all rules are enabled. Instead, tune your rule set to match your network environment as closely as possible.

Set intrusion rule action

Set intrusion rule actions to control how the system responds to detected threats, allowing you to configure rules to alert on or block specific intrusion attempts.

Intrusion rule actions are policy-specific.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Click **Snort 3 Version** next to the policy you want to edit.

Tip
This page shows the total number of:

- disabled rules
- enabled rules set to Alert
- enabled rules set to Block
- overridden rules

Step 3 Choose the rule or rules where you want to set the rule action.

Step 4 Choose one of the rule actions from the **Rule Action** drop-down list. See Edit Snort 3 Intrusion Policies, on page 2821 for more information about the different rule actions.

Step 5 Click **Save**.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Intrusion event notification filters in an intrusion policy

An intrusion event notification filter in an intrusion policy is a security mechanism that

- determines event importance based on frequency of occurrence or source and destination IP addresses,
- allows administrators to specify thresholds for when events require attention, and
- prevents system overwhelm by limiting notifications for widespread attacks.

Event notification filtering scenarios

The importance of an intrusion event can be based on frequency of occurrence, or on source or destination IP address. In some cases you may not care about an event until it has occurred a certain number of times. For example, you may not be concerned if someone attempts to log into a server until they fail a certain number of times. In other cases, you may only need to see a few occurrences to know there is a widespread problem. For example, if a DoS attack is launched against your web server, you may only need to see a few occurrences of an intrusion event to know that you need to address the situation. Your system is overwhelmed when it receives hundreds of the same event.

Intrusion event thresholds

An intrusion event threshold is a configuration setting that

- limits the number of times the system logs and displays an intrusion event based on how many times the event is generated within a specified time period,
- prevents you from being overwhelmed with a large number of identical events, and
- can be set per shared object rule, standard text rule, or inspector rule.

Set intrusion event thresholds

To set a threshold, first specify the thresholding type.

Table 153: Thresholding options

Option	Description
Limit	Logs and displays events for the specified number of packets (specified by the Count argument) that trigger the rule during the specified time period. For example, if you set the type to **Limit**, the **Count** to `10`, and the **Seconds** to `60`, and 14 packets trigger the rule, the system stops logging events for the rule after displaying the first 10 that occur within the same minute.
Threshold	Logs and displays a single event when the specified number of packets (specified by the Count argument) trigger the rule during the specified time period. Note that the counter for the time restarts after you hit the threshold count of events and the system logs that event. For example, you set the type to **Threshold**, **Count** to `10`, and **Seconds** to `60`, and the rule triggers 10 times by second 33. The system generates one event, then resets the Seconds and Count counters to 0. The rule then triggers another 10 times in the next 25 seconds. Because the counters reset to 0 at second 33, the system logs another event.
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to **Both**, **Count** to two, and **Seconds** to `10`, these event counts result: • If the rule is triggered once in 10 seconds, the system does not generate any events (the threshold is not met) • If the rule is triggered twice in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time) • If the rule is triggered four times in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time, and subsequent events are ignored)

Secondly, specify tracking, which determines whether the event threshold is calculated per source or destination IP address.

Table 154: Thresholding IP options

Option	Description
Source	Calculates event instance count per source IP address.
Destination	Calculates event instance count per destination IP address.

Finally, specify the number of instances and time period that define the threshold.

Table 155: Thresholding Instance/Time options

Option	Description
Count	The number of event instances per specified time period per tracking IP address required to meet the threshold.

Option	Description
Seconds	The number of seconds that elapse before the count resets. If you set the threshold type to **limit**, the tracking to **Source IP**, the **count** to `10`, and the **seconds** to `10`, the system logs and displays the first 10 events that occur in 10 seconds from a given source port. If only 7 events occur in the first 10 seconds, the system logs and displays those; if 40 events occur in the first 10 seconds, the system logs and displays 10, then begins counting again when the 10-second time period elapses.

You can use intrusion event thresholding alone or combine it with rate-based attack prevention, the `detection_filter` keyword, or intrusion event suppression.

Tip You can also add thresholds from within the packet view of an intrusion event.

Set threshold for an intrusion rule in snort 3

Set a threshold for an intrusion rule to control when and how often alerts are generated based on event instances within a specified time period.

You can set a single threshold for a rule from the Rule Detail window. Adding a threshold overwrites any existing threshold for the rule. The threshold you set for an intrusion rule is applied to each packet thread. However, the configuration is fully applied only within the context of a unique flow. You might see more alerts on different network flows, but you will not see fewer alerts than the configured number.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 From an intrusion rule's Alert Configuration column, click the **None** link.

Step 4 Click **Edit** ( ).

Step 5 In the Alert Configuration window, click the **Threshold** tab.

Step 6 From the **Type** drop-down list, choose the type of threshold you want to set:

- Choose **Limit** to limit notification to the specified number of event instances per time period.
- Choose **Threshold** to provide notification for each specified number of event instances per time period.
- Choose **Both** to provide notification once per time period after a specified number of event instances.

Step 7 Choose **Source** or **Destination** in the **Track By** field to indicate whether you want the event instances tracked by source or destination IP address.

Step 8 Enter the number of event instances you want to use as your threshold in the **Count** field.

Step 9 Enter a number that specifies the time period, in seconds, for which event instances are tracked in the **Seconds** field.

Step 10 Click **Save**.

Refer to the video Snort 3 Suppression and Threshold for additional support and information.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 202.

View and delete intrusion event thresholds

This task allows you to view existing threshold settings for intrusion rules and delete them when they are no longer appropriate for your system configuration.

To view or delete an existing threshold setting for a rule, use the Rules Details view to display the configured settings for a threshold and see if they are appropriate for your system. If they are not, you can add a new threshold to overwrite the existing values.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Choose the rule with a configured threshold as shown in the **Alert Configuration** column (the **Alert Configuration** column displays **Threshold** as a link for the rule).

Step 4 To remove the threshold for the rule, click **Threshold** link in the **Alert Configuration** column.

Step 5 Click **Edit** ( ).

Step 6 Click **Threshold** tab.

Step 7 Click **Reset**.

Step 8 Click **Save**.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Intrusion policy suppression configuration

Intrusion policy suppression configuration is a security mechanism that

- suppresses intrusion event notification when a specific IP address or range of IP addresses triggers a specific rule or inspector,
- eliminates false positives by allowing rules to trigger for all packets while only showing events for legitimate attacks, and
- is useful for scenarios such as mail servers that transmit packets resembling specific exploits.

Mail server suppression example

If you have a mail server that transmits packets that look like a specific exploit, you might suppress event notification for that event when it is triggered by your mail server. The rule triggers for all packets, but you only see events for legitimate attacks.

Intrusion policy suppression types

You can use intrusion event suppression alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, or intrusion event thresholding.

Tip You can add suppressions from within the packet view of an intrusion event. You can also access suppression settings by using the **Alert Configuration** column on the intrusion rules editor page (**Objects** > **Intrusion Rules**, click **Snort 3 All Rules**).

Set suppression for an intrusion rule in Snort 3

This task allows you to configure one or more suppressions for a rule in your intrusion policy to reduce false positives and focus on relevant security events.

You can set one or more suppressions for a rule in your intrusion policy.

Before you begin

Ensure you create the required network objects to be added for source or destination suppression.

Follow these steps to set suppression for an intrusion rule in Snort 3:

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Click the **None** link in the Alert Configuration column for the intrusion rule.

Step 4 Click **Edit** ( ).

Step 5 From the **Suppressions** tab, click the add icon **Add** ( ) next to any of these options:

- Choose **Source Networks** to suppress events generated by packets originating from a specified source IP address.
- Choose **Destination Networks** to suppress events generated by packets going to a specified destination IP address.

Step 6 Select any of the preset networks in the **Network** drop-down list.

Step 7 Click **Save**.

Step 8 (Optional) Repeat steps 5, 6, and 7 if needed.

Step 9 Click **Save** in the Alert Configuration window.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

View and delete suppression conditions

View or delete existing suppression conditions to manage event notifications for specific network conditions.

You may want to view or delete an existing suppression condition. For example, you can suppress event notification for packets originating from a mail server IP address because the mail server normally transmits packets that look like exploits. If you then decommission that mail server and reassign the IP address to another host, you should delete the suppression conditions for that source IP address.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Choose the rule for which you want to view or delete suppressions.

Step 4 Click **Suppression** in the **Alert Configuration** column.

Step 5 Click **Edit** ( ).

Step 6 Click **Suppressions** tab.

Step 7 Remove the suppression by clicking **Clear** ( ) next to the suppression.

Step 8 Click **Save**.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Add intrusion rule comments

Add comments to rules in your intrusion policy for documentation and management purposes.

You can add comments to rules in your intrusion policy. Comments added this way are specific to the policy. Comments you add to a rule in one intrusion policy are not visible in other intrusion policies.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Click **Snort 3 Version** next to the policy you want to edit.

Step 3 In the right side of the page where all the rules are listed, choose the rule where you want to add a comment.

Step 4 Click **Comment** ( ) under the **Comments** column.

Step 5 In the **Comments** field, enter the rule comment.

Step 6 Click **Add Comment**.

Step 7 Click **Save**.

The system displays a **Comment** ( ) next to the rule in the Comments column.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Prepare custom rules for snort 3 conversion

Before converting from Snort 2 to Snort 3, make sure you are prepared to manage your custom rule set for Snort 3.

- If you are using a rule set from a third-party vendor, contact that vendor to confirm that their rules will successfully convert to Snort 3 or to obtain a replacement rule set written natively for Snort 3.
- If you have custom rules that you have written yourself, familiarize yourself with writing Snort 3 rules before converting, so you can update your rules to optimize Snort 3 detection after conversion.

Learn more about writing rules in Snort 3:

- https://blog.snort.org/2020/08/how-rules-are-improving-in-snort-3.html
- https://blog.snort.org/2020/10/talos-transition-to-snort-3.html

You can refer to other blogs at https://blog.snort.org/ to learn more about Snort 3 rules.

Important Snort 2 network analysis policy (NAP) settings *cannot* be copied to Snort 3 automatically. NAP settings must be manually replicated in Snort 3.

Convert all Snort 2 custom rules across all intrusion policies to Snort 3

Convert legacy Snort 2 custom rules to the updated Snort 3 format to ensure compatibility with current intrusion detection systems.

This task migrates existing Snort 2 custom rules to Snort 3 format, allowing you to maintain your custom rule configurations while upgrading to the latest intrusion detection engine. You can either import the converted rules directly or download them for review before importing.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Ensure **All Rules** is selected in the left pane.

Step 4 Click the **Tasks** drop-down list and choose:

- **Convert Snort 2 rules and import**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and import them into Firewall Management Center as Snort 3 custom rules.
- **Convert Snort 2 rules and download**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and download them into your local system.

Step 5 Click **OK**.

Note

- If you selected **Convert and import** in the previous step, then all the converted rules are saved under a newly created rule group **All Snort 2 Converted Global** under **Local Rules**.
- If you selected **Convert and download** in the previous step, then save the rules file locally. You can review the converted rules in the downloaded file and later upload them by following the steps in Add custom rules to rule groups, on page 2846.

Refer to the video Converting Snort 2 Rules to Snort 3 for additional support and information.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Convert Snort 2 custom rules of a single intrusion policy to Snort 3

Convert Snort 2 custom rules in an intrusion policy to Snort 3 format to maintain compatibility with Snort 3 environments.

When migrating to Snort 3, existing Snort 2 custom rules need to be converted to maintain their functionality. This task allows you to convert custom rules for a specific intrusion policy.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 In the **Intrusion Policies** tab, click **Show Snort 3 Sync status**.

Step 3 Click the **Sync** icon **Snort out-of-Sync** ( ) of the intrusion policy.

Note

If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in green**Snort in-Sync** ( ). It indicates that there are no custom rules to be converted.

Step 4 Read the summary and click the **Custom Rules** tab.

Step 5 Choose:

- **Import converted rules to this policy**: To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and import them into Firewall Management Center as Snort 3 custom rules.

- **Download converted rules**: To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and download them into your local system. You can review the converted rules in the downloaded file and later upload the file by clicking the upload icon.

Step 6 Click **Re-Sync**.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Add custom rules to rule groups

Upload custom rules to add locally created Snort 3 rules to the list of all available rules in the system.

Uploading custom rules in the Firewall Management Center adds the custom rules that you have created locally to the list of all the Snort 3 rules.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Click the **Tasks** drop-down list.

Step 4 Click **Upload Snort 3 Rules**.

Step 5 Drag and drop the `.txt` or `.rules` file that contains the Snort 3 custom rules that you have created.

Step 6 Click **OK**.

Note
If there are any errors in the selected file, then you cannot proceed further. You can download the error file and **Replace File** link to upload version 2 of the file, after fixing the errors.

Step 7 Associate rules to a rule group to add the new rules to that group.

You can also create a new custom rule group (by clicking the **Create New Custom Rule Group** link) and then add the rules to the new group.

Note
If there are no existing local rule groups, then proceed by clicking **Create New Custom Rule Group to proceed**. Enter a **Name** for the new rule group and click **Save**.

Step 8 Choose either of these options:

- **Merge Rules** to merge the new rules that you are adding with the existing rules in the rule group.
- **Replace all rules in the group with file contents** to replace all the existing rules with the new rules that you are adding.

Note
If you chose more than one rule group in the previous step, then only the **Merge Rules** option is available.

Step 9 Click **Next**.

Review the summary to know the new rule IDs that are being added and optionally download it.

Step 10 Click **Finish**.

Important All uploaded rules are disabled by default. Change their state as needed to ensure the rules are active.

What to do next

- Uploading custom rules in the Firewall Management Center adds the custom rules that you have created to the list of all the Snort 3 rules. To enforce these custom rules on the traffic, add and enable these rules in the required intrusion policies. For information on adding rule groups with custom rules to an intrusion policy, see Add rule groups with custom rules to an intrusion policy, on page 2847. For information on enabling custom rules, see Manage custom rules in Snort 3, on page 2847.
- Deploy configuration changes. See Deploy configuration changes, on page 2806.

Add rule groups with custom rules to an intrusion policy

Enable custom rules in an intrusion policy to enforce these rules on network traffic.

Custom rules that are uploaded in the system have to be enabled in an intrusion policy to enforce those rules on the traffic. After uploading custom rules on Firewall Management Center, add the rule group with the new custom rules in the intrusion policy.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 In the **Intrusion Policies** tab, click the **Snort 3 Version** of the intrusion policy.

Step 3 Click **Add** (+) next to the Rule Groups search bar.

Step 4 In the **Add Rule Groups** window, click the **Expand Arrow** (›) icon next to a rule group to expand the local rule group.

Step 5 Check the check box next to the uploaded custom rules group.

Step 6 Click **Save**.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Manage custom rules in Snort 3

This task enables uploaded custom rules in intrusion policies to enforce rule monitoring on network traffic.

Custom rules that are uploaded in the system have to be added to an intrusion policy and enabled to enforce those rules on the traffic. You can enable the uploaded custom rules across all policies or selectively on individual policies.

Before you begin

Follow these steps to enable custom rules in one or many intrusion policies:

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Expand **Local Rules**.

Step 4 Select the required rule group.

Step 5 Select the rules by checking the check boxes next to them.

Step 6 Select **Per Intrusion Policy** from the **Rule Actions** drop-down list.

Step 7 Choose:

- **All Policies** to have the same rule actions for all the rules to be added.
- **Per Intrusion Policy** to have different rule actions for each intrusion policy.

Step 8 Set the rule actions:

- If you selected All Policies in the previous step, then select the required rule action from the **Select Override state** drop-down list.
- If you selected Per Intrusion Policy in the previous step, then select the **Rule Action** against the policy name. To add more policies, click **Add Another**.

Step 9 Optionally, add a comment in the **Comments** text box.

Step 10 Click **Save**.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Delete custom rules

Delete custom rules that are no longer needed from your intrusion rule configuration to maintain an organized and efficient rule set.

Use this procedure when you need to remove custom rules from the Local Rules section in the Snort 3 All Rules interface. Before deletion, all rules must be disabled to prevent configuration conflicts.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Expand **Local Rules** in the left pane.

Step 4 Check the check boxes of the rules you want to delete.

Step 5 Ensure that the rule action for all the rules that you select is **Disable**.

If required, follow these steps to disable the rule action for multiple selected rules:

a) From the **Rule Actions** drop-down box, select **Per Intrusion Policy**.
b) Select **All Policies** radio button.
c) Select **Disable** from the **Select Override state** drop-down list.
d) Click **Save**.
e) Check the check boxes of the rules you want to delete.

Step 6 From the **Rule Actions** drop-down list, select **Delete**.

Step 7 Click **Delete** in the Delete Rules pop-up window.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Delete rule groups

Delete rule groups to remove unwanted or obsolete intrusion detection rules from your system configuration and maintain an organized rule structure.

Before deleting a rule group, you must exclude it from all intrusion policies and disable all rules within the group to prevent system conflicts.

Before you begin

Exclude the rule group you want to delete from all intrusion policies where you have included it. For steps on excluding a rule group from an intrusion policy, see Edit Snort 3 Intrusion Policies, on page 2821.

Follow these steps to delete rule groups:

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Expand **Local Rules** in the left pane.

Step 4 Select the rule group to be deleted.

Step 5 Ensure the rule action for all the rules in the group is set to **Disable** before proceeding.

If the rule action for any of the rules is anything other than **Disable**, then you cannot delete the rule group. If required, follow the steps below to disable the rule action for all the rules:

a) Check the check box below the **Rule Actions** drop-down list to select all the rules in the group.
b) From the **Rule Actions** drop-down box, select **Per Intrusion Policy**.
c) Select **All Policies** radio button.
d) Select **Disable** from the **Select Override state** drop-down list.
e) Click **Save**.

Step 6 Click the **Delete** ( ) next to the rule group.

Step 7 Click **OK** in the Delete Rule Group pop-up window.

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Recommended rules

This section provides an insight into Secure Firewall recommended rules and generating and applying Secure Firewall recommended rules.

Snort 3 rule changes in LSP updates

Snort 3 rule changes in LSP updates are modifications to system-defined intrusion rules that

- replace existing rules with new rules for better detection capabilities,
- combine multiple rules into a single rule or expand a single rule into multiple rules when better detection is possible, and
- remove existing system-defined rules for better management as part of the LSP update.

To get notifications for changes to any *overridden* system-defined rules during LSP updates, ensure that the **Retain user overrides for deleted Snort 3 rules** check box is checked.

To navigate to the **Retain user overrides for deleted Snort 3 rules** check box, click **System** ( ) > **Configuration** > **Intrusion Policy Preferences**.

By default this check box is checked. When this check box is checked, the system retains the rule overrides in the new replacement rules that are added as a part of the LSP update. The notifications are shown in the **Tasks** tab under the Notifications icon that is located next to **System** ( ).

Secure Firewall recommended rules

Secure Firewall recommended rules are intrusion rule recommendations that

- target vulnerabilities associated with host assets detected in the network, such as operating systems, servers, and client application protocols,
- allow you to tailor your intrusion policy to the specific needs of your monitored network, and

- provide individual recommendations for each intrusion policy. These typically apply to standard text rules and shared object rules, but recommendations can also include inspector and decoder rules.

Rule recommendation configuration options

When you generate rule state recommendations, you can use the default settings or configure advanced settings. Advanced settings allow you to:

- Redefine which hosts on your network the system monitors for vulnerabilities
- Influence which rules the system recommends based on rule overhead
- Specify whether to generate recommendations to disable rules

You can also choose to use the recommendations immediately or review the recommendations (and affected rules) before accepting them.

Choosing to use recommended rule states adds a read-only Secure Firewall Recommendations layer to your intrusion policy, and subsequently choosing not to use recommended rule states removes the layer.

You can schedule a task to generate recommendations automatically based on the most recently saved configuration settings in your intrusion policy.

The system does not change rule states that you set manually such as:

- Manually setting the states of specified rules *before* you generate recommendations prevents the system from modifying the states of those rules in the future.
- Manually setting the states of specified rules *after* you generate recommendations overrides the recommended states of those rules.

Tip The intrusion policy report can include a list of rules with rule states that differ from the recommended state.

While displaying the recommendation-filtered Rules page, or after accessing the Rules page directly from the navigation panel or the Policy Information page, you can manually set rule states, sort rules, and take any of the other actions available on the Rules page, such as suppressing rules, setting rule thresholds, and so on.

Note The Cisco Talos Intelligence Group (Talos) determines the appropriate state of each rule in the system-provided policies. If you use a system-provided policy as your base policy, and you allow the system to set your rules to the Secure Firewall recommended rule state, the rules in your intrusion policy match the settings recommended for your network assets.

Generate new Secure Firewall recommendations in Snort 3

Generate Secure Firewall recommendations to create new recommended rule settings in Snort 3.

Generate the Secure Firewall recommendations for the intrusion policy and then follow the steps that are listed here to create new recommended rule settings in Snort 3. Rule overheads are interpreted as **security levels**

based on the threshold policies you select in Snort 3. The recommended action depends on the selected security level. If it is higher than the base policy, the recommendation includes actions beyond only generating events.

Prior to setting the Secure Firewall recommendations you should ask which of the three points listed below closely matches the goal:

- Increased Protection —Enable additional rules based on vulnerabilities found in the host database and do not automatically disable any rules. This will likely result in a larger rule set.
- Focused Protection—Enable additional rules and disable existing rules based on vulnerabilities found in the host database. This can increase or decrease the number of rules depending on vulnerabilities discovered.
- Higher Efficiency—Use the currently enabled rule set and disable any rules for vulnerabilities not found in the host database. This will likely result in a smaller enabled rule set.

Based on the response, these are the recommended actions:

- Set recommendations to the next highest security level, and clear the option to disable rules.
- Set recommendations to the next highest security level, and select the option to disable rules.
- Set recommendations to the current security level, and select the option to disable rules.

Before you begin

Secure Firewall recommendations have the following requirements:

- Ensure that hosts are present in the system to generate recommendations.
- Protected networks configured for recommendations should map to the hosts present in the system

Follow these steps to generate new Secure Firewall recommendations in Snort 3:

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Click the **Snort 3 Version** button of the intrusion policy.

Step 3 Click the **Recommendations (Not in Use)** layer to configure the rule recommendations. Click **Start**.

In the Secure Firewall Rule Recommendations window you can set the following:

- **Security Level**: Click to select the security level. Optionally, you can check the **Accept Recommendation to Disable Rules** checkbox to disable rules that are not enabled at the input security level and in protected networks. Only enable this option if you need to trim your rule set due to a high number of alerts or to improve inspection performance. The security levels are:
  - Security level 1: Connectivity Over Security

    No Impact—No new rules are enabled and no existing rules are disabled. To increase the protection, select a higher security level.

    Lower Security (checkbox is checked)—All rules are disabled except for the rules in the Connectivity Over Security ruleset that match potential vulnerabilities on discovered hosts. It is recommended instead to adjust the Base Policy.

- Security level 2: Balanced Security Over Connectivity

  No Impact—No new rules are enabled and no existing rules are disabled. To increase the protection, select a higher security level.

  Higher Efficiency(checkbox is checked)—Keeps existing rules that match potential vulnerabilities on discovered hosts and disables rules for vulnerabilities not found on the network.

- Security level 3: Security Over Connectivity

  Increased Security—Enables additional rules that match potential vulnerabilities on discovered hosts based on the Maximum Detection ruleset.

  Focused Security(checkbox is checked)—Enables additional rules that match vulnerabilities on discovered hosts based on the Security Over Connectivity ruleset, while disabling existing rules that do not match potential vulnerabilities on discovered hosts.

- Security level 4 : Maximum Detection

  Increased Security—Enables additional rules that match potential vulnerabilities on discovered hosts based on the Security Over Connectivity ruleset.

  Focused Security (checkbox is checked)—Enables additional rules that match vulnerabilities on discovered hosts based on the Maximum Detection ruleset, while disabling existing rules that do not match potential vulnerabilities on discovered hosts.

  Note
  Maximum Detection enables a very high number of rules and may impact performance. We recommend you to review and test this setting before deploying into a production environment.

- **Protected Networks**: Specifies the monitored networks or individual hosts to examine for recommendations. You can select one or more system or custom defined network objects from the drop-down list. By default, any IPv4 or IPv6 networks are selected, if no selection is done.

  Important
  The Secure Firewall Rule Recommendations depend on network discovery. Protected Networks apply to any hosts discovered within the ranges configured in your Network Discovery policy. For more information, see the chapter Network Discovery Policies in the *Cisco Secure Firewall Management Center Device Configuration guide*.

  Click the **Add +** button to create a new network object of type Host or Network and click **Save**.

Step 4 Generate and apply recommendations:

- **Generate**: Generates the recommendations for an intrusion policy. This action lists the rules under Recommended Rules (Not in use).
- **Generate and Apply**: Generates and applies the recommendations for an intrusion policy. This action lists the rules under Recommended Rules (In use).

Recommendations are generated successfully. A new recommendation tab appears with all the recommended rules with their corresponding recommended actions. Rule action preset filters are also available for this tab, in addition with new recommendations.

Step 5 You can verify the recommendations and choose to apply them.

- **Accept**—Applies the previously generated recommendations for an Intrusion policy.

- **Refresh**—Regenerates and updates the rule recommendations for an Intrusion policy.
- **Edit**—It opens the Recommendations dialog box, you can provide the recommendation input values and then generate the recommendations.
- **Remove All**—Revert or remove the applied recommended rules from the policy and also removes the recommendation tab.

Under **All Rules**, there is a Recommended Rules section which displays the recommended rules.

Note
Final action for an Intrusion rule is applied based on the rule action priority order and following is the rule action priority order:

Rule Override > Generated Recommendations > Group Override > Base Policy Default Action

For enabled recommendations, Firewall Management Center considers the current state: group overrides, base policy, and recommendation configurations and priority order of actions is:

pass > block > reject > drop > rewrite > alert

What to do next

Deploy configuration changes. See Deploy configuration changes, on page 2806.

Mitigate threats using MITRE framework in Snort 3 intrusion policies

This section describes how to use the MITRE ATT&CK framework in Snort 3 intrusion policies to identify attack tactics and techniques, review rule coverage, and improve threat mitigation.

MITRE framework

The MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK) framework is a comprehensive knowledge base that

- provides insights into the tactics, techniques, and procedures (TTPs) distributed by threat actors aiming to harm systems,
- incorporates groups to refer to threat groups, activity groups, or threat actors based on the set of tactics and techniques they employ, and
- is compiled as matrices, with each matrix representing an operating system or a particular platform.

Each technique includes information about execution, procedures, defenses, detections, and real-world examples.

Note See https://attack.MITRE.org for additional information about MITRE ATT&CK.

The management center uses the MITRE ATT&CK Framework to enhance threat detection and response, incorporating these capabilities:

- Intrusion events include TTPs, allowing administrators to manage traffic with greater granularity by grouping rules according to vulnerability type, target system, or threat category.
- Select malware events use TTPs, enhancing the ability to detect and respond to threats.

MITRE framework benefits

MITRE framework provides security advantages that

- enable administrators to act on traffic based on MITRE Tactics, Techniques, and Procedures (TTPs) added to intrusion events,
- allow viewing and handling traffic with more granularity through grouping rules by vulnerability type, target system, or threat category, and
- provide the ability to organize intrusion rules according to the MITRE ATT&CK framework for customized policies based on specific attacker tactics and techniques.

Sample business scenario for MITRE network

A sample business scenario for MITRE network is a security implementation example that

- uses Snort 3 as the primary intrusion detection and prevention system for a large corporate network,
- addresses the need for robust network security measures in a rapidly evolving threat landscape, and
- enables network administrators to determine if configured policies are finding traffic of interest and tracking known attack groups.

Network security considerations

Network administrators need to know if adversaries are attempting to take advantage of weaknesses in their systems or applications to cause unanticipated behavior. The weakness in the system can be a bug, a glitch, or a design vulnerability. The applications may be websites, databases, standard services, such as Server Message Block (SMB) or Secure Shell (SSH), network device administration and management protocols or applications, such as web servers and related services.

The insights provided by the MITRE framework provide administrators with a more precise opportunity to specify protection for specific assets and protect their network from specific threat groups.

Prerequisites for MITRE framework

Ensure that your system meets these requirements before implementing the MITRE framework:

- You must be running Secure Firewall Management Center and Secure Firewall Threat Defense Version 7.3.0 or later with Snort 3.
- You must have at least one intrusion policy. See Create a custom snort 3 intrusion policy, on page 2820.

View and edit your Snort 3 intrusion policy

This task allows you to customize your Snort 3 intrusion policy by modifying rule groups, security levels, and accessing Cisco-recommended rules to enhance threat detection capabilities.

Snort 3 intrusion policies use a hierarchical rule group structure organized by the MITRE framework. You can drill down to specific rule categories and modify security levels for bulk changes or individual rule actions. The policy editor provides multiple layers, including Group Overrides, Rule Overrides, Recommendations, and Summary views.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

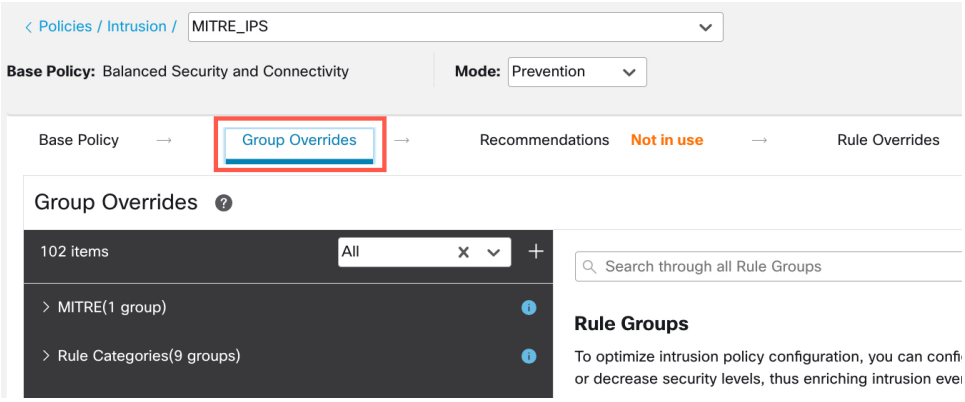
Step 2 Ensure that the **Intrusion Policies** tab is chosen.

a) Click **Snort 3 Version** next to the intrusion policy that you want to view or edit.

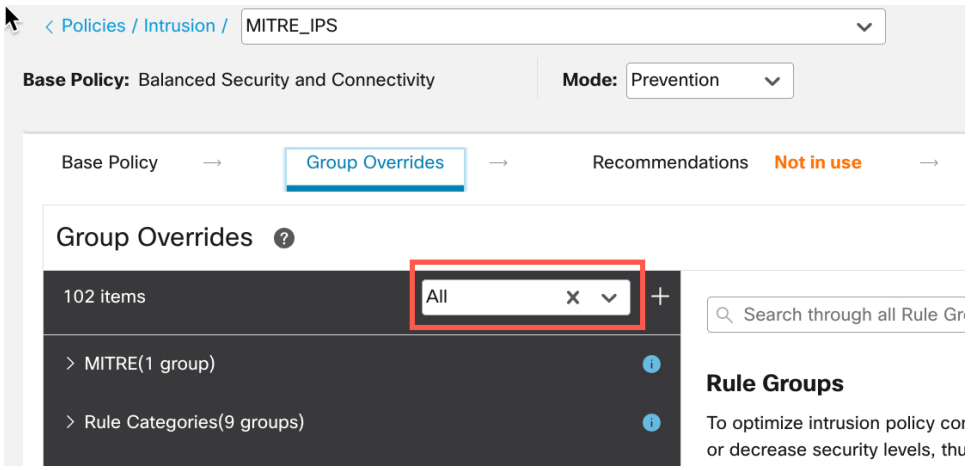
b) Close the Snort helper guide that is displayed.

Step 3 Click the **Group Overrides** layer.

This layer lists all the categories of rule groups, in a hierarchical structure. You can drill down to the final leaf rule group under each rule group.

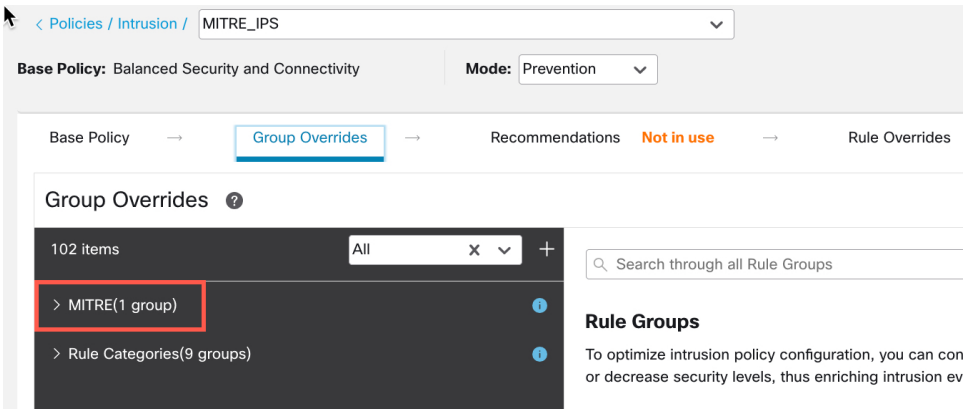


a) Under **Group Overrides**, ensure that **All** is chosen in the drop-down list, so that all the rule groups for the corresponding intrusion policy are visible in the left pane.

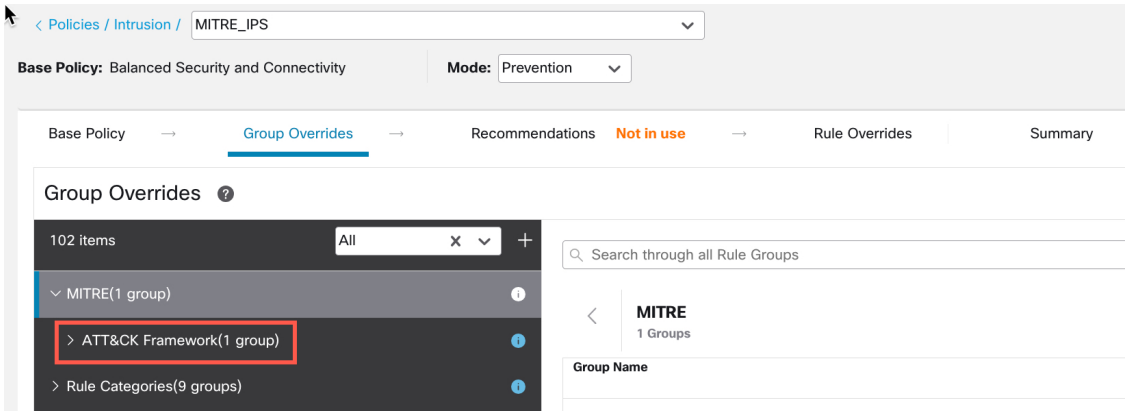


Step 4 Click **MITRE** in the left pane.

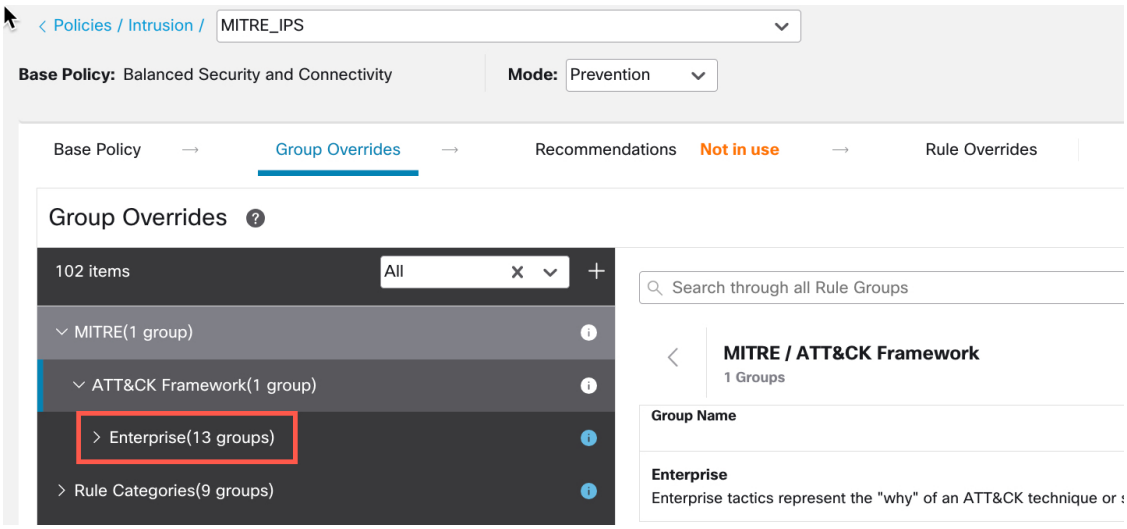
Note
Depending on your specific requirements, you can choose the **Rule Categories** rule group or any other rule group and subrule groups under it. All the rule groups use the MITRE framework.



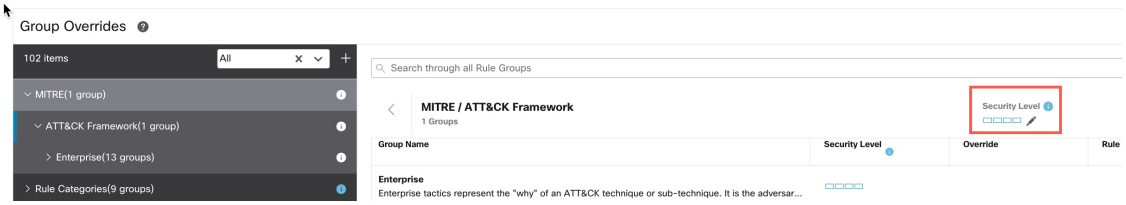
a) Under **MITRE**, click **ATT&CK Framework** to drill down.



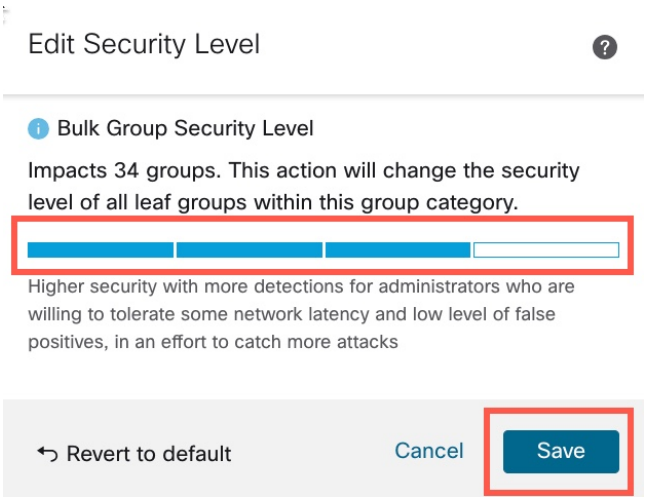
b) Under **ATT&CK Framework**, click **Enterprise** to expand it.



c) Click the **Edit** (✎) icon next to the **Security Level** of the rule group to make bulk changes to all associated rule group security levels under the **Enterprise** rule group category.

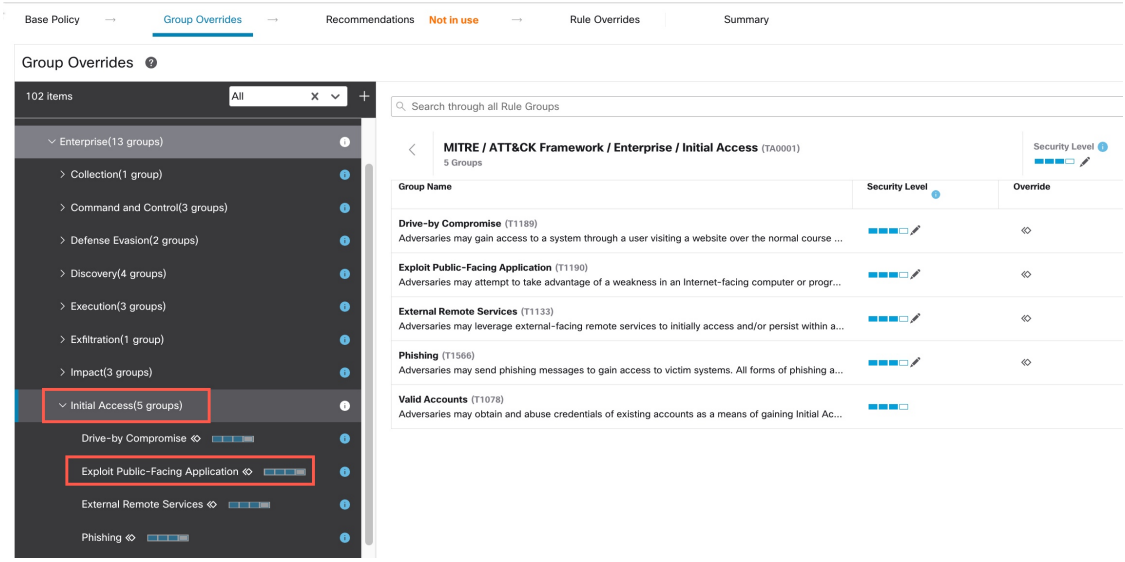


d) In the **Edit Security Level** window, choose a **Security Level** (in this example, **3**), and click **Save**.

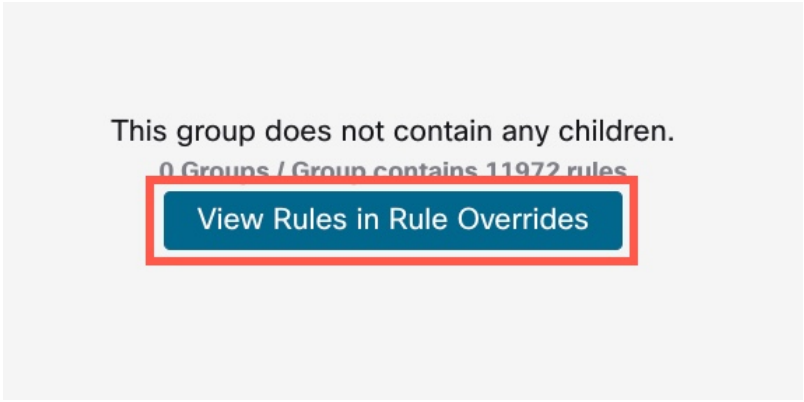


Step 5 Under **Enterprise**, click **Initial Access** to expand it.

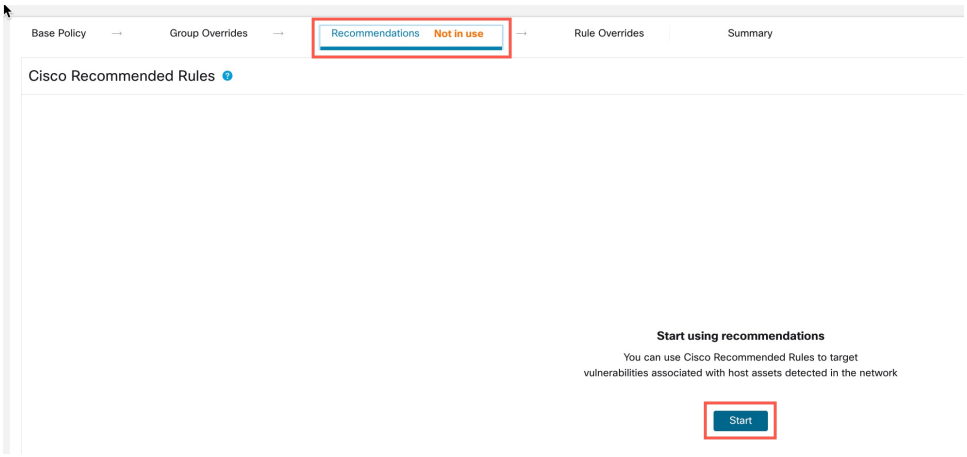
Step 6 Under **Initial Access**, click **Exploit Public-Facing Application**, which is the last leaf group.



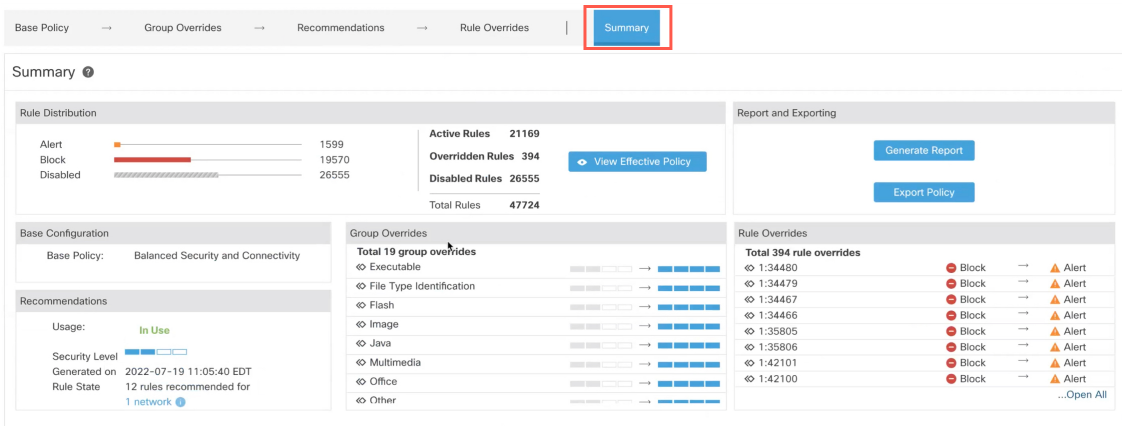
Step 7 Click **View Rules in Rule Overrides** to view the available rules, their details, and actions. You can change the rule actions for one or multiple rules in the **Rule Overrides** layer.



Step 8 Click the **Recommendations** layer and then click **Start** to start using Cisco-recommended rules. You can use the intrusion rule recommendations to target the vulnerabilities that are associated with the host assets detected in the network. For more information, see Generate new Secure Firewall recommendations in Snort 3, on page 2851.



Step 9 Click the **Summary** layer for a holistic view of the current changes to the policy. Based on the rule overrides, security-level changes, and generation of Cisco-recommended rules, you can view the rule distribution of the policy, group overrides, rule overrides, rule recommendations, and additional details to verify your changes.



What to do next

Deploy your intrusion policy to detect and log events that are triggered by the Snort rules. See Deploy configuration changes, on page 2806.

View intrusion events

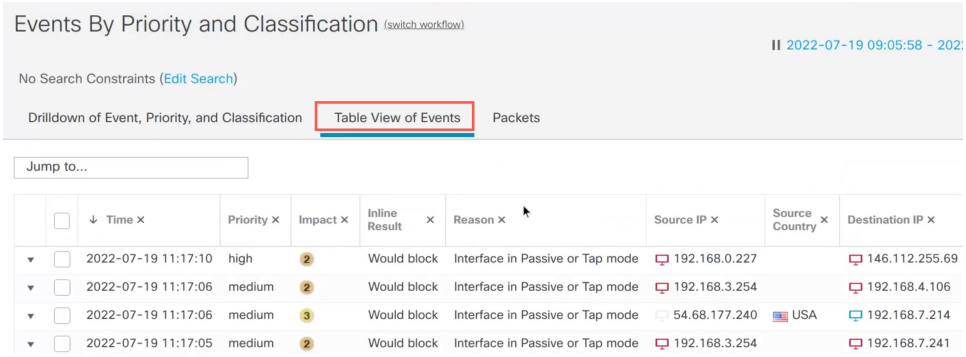
This task allows you to view MITRE ATT&CK techniques and rule groups associated with intrusion events, providing security context and attack pattern mapping for better threat analysis.

You can view the MITRE ATT&CK techniques and rule groups in the intrusion events on the **Classic Event Viewer** and **Unified Event Viewer** pages. Talos provides mappings from Snort rules (GID:SID) to MITRE ATT&CK techniques and rule groups. These mappings are installed as part of the Lightweight Security Package (LSP).

Procedure

Step 1 Click **Analysis** and select **Events** under **Intrusions**.

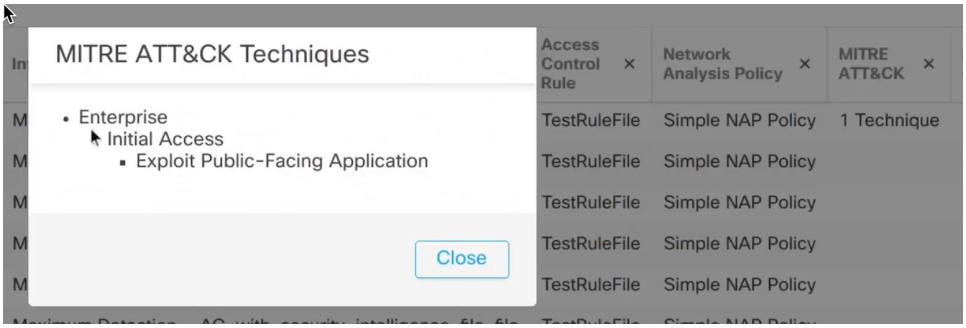
Step 2 Click the **Table View of Events** tab.



Step 3 Under **MITRE ATT&CK**, you can see the techniques for an intrusion event. Click **1 Technique** to view the MITRE ATT&CK techniques.



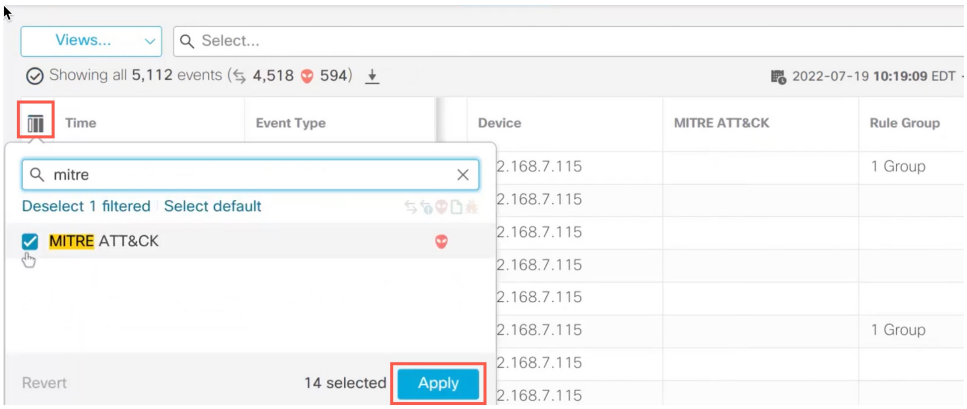
In this example, **Exploit Public-Facing Application** is the technique.



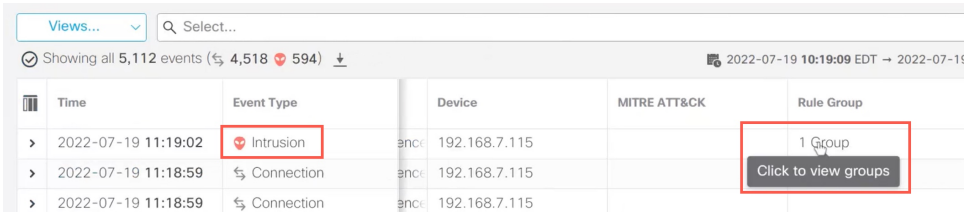
Step 4 Click **Close**.

Step 5 Click **Analysis** and select **Unified Events**.

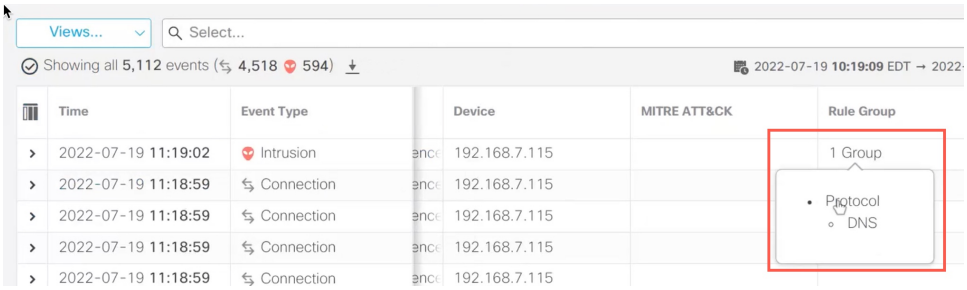
Step 6 If not enabled, click the column selector icon to enable the **MITRE ATT&CK** and **Rule Group** columns.



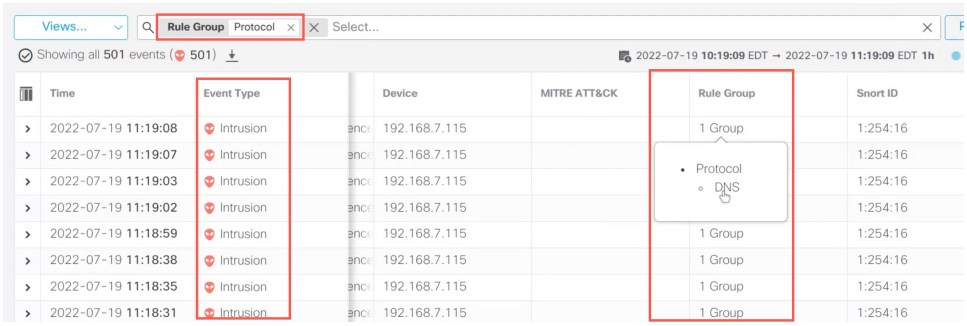
Step 7 In this example, the intrusion event is triggered by an event that is mapped to one rule group. Click **1 Group** under the **Rule Group** column.



Step 8 You can view **Protocol**, which is the parent rule group, and the DNS rule group under it. Choose **Protocol** > **DNS** to search for all the intrusion events that have at least one rule group.



The search results are displayed.



Additional references

This reference provides links to additional documentation and resources for intrusion policies in Snort 3 and MITRE framework information for malware event classification.

- Intrusion Policy in Snort 3
- Edit Snort 3 Intrusion Policies, on page 2821
- MITRE Information in Malware Events



CHAPTER 100

Intrusion Prevention Performance Tuning

This chapter explains how to refine intrusion prevention performance.

- Intrusion prevention performance tuning, on page 2865
- License requirements for intrusion prevention performance tuning, on page 2866
- Requirements and prerequisites for intrusion prevention performance tuning, on page 2866
- Limiting pattern matching for intrusions, on page 2866
- Regular expression limits overrides for intrusion rules, on page 2867
- Override regular expression limits for intrusion rules, on page 2868
- Per packet intrusion event generation limits, on page 2869
- Limiting intrusion events generated per packet, on page 2869
- Packet and intrusion rule latency threshold configuration, on page 2870
- Intrusion performance statistic logging configuration, on page 2877
- Configure intrusion performance statistic logging, on page 2878

Intrusion prevention performance tuning

Intrusion prevention performance tuning is a system optimization mechanism that

- improves performance while analyzing traffic for attempted intrusions,
- applies settings on a per-access control policy basis to all intrusion policies invoked by that parent access control policy, and
- balances security with the need to maintain device latency at an acceptable level.

Performance tuning features

Cisco provides several features for improving the performance of your system as it analyzes traffic for attempted intrusions. You can:

- specify the number of packets to allow in the event queue. You can also, before and after stream reassembly, enable or disable inspection of packets that will be rebuilt into larger streams.
- override default match and recursion limits on PCRE that are used in intrusion rules to examine packet payload content.
- elect to have the rules engine log more than one event per packet or packet stream when multiple events are generated, allowing you to collect information beyond the reported event.

- balance security with the need to maintain device latency at an acceptable level with packet and rule latency thresholding.
- configure the basic parameters of how devices monitor and report their own performance. This allows you to specify the intervals at which the system updates performance statistics on your devices.

License requirements for intrusion prevention performance tuning

This reference provides licensing information required for implementing intrusion prevention performance tuning on Cisco Secure Firewall devices, covering both Threat Defense and Classic licensing models.

Threat Defense License

IPS

Classic License

Protection

Requirements and prerequisites for intrusion prevention performance tuning

This reference provides the system requirements and prerequisite configurations needed before implementing intrusion prevention performance tuning on Cisco Secure Firewall devices.

Model support

Any

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Limiting pattern matching for intrusions

Limit pattern matching for intrusions to optimize system performance by controlling how many pattern states are analyzed per packet and managing content inspection before and after stream reassembly.

Pattern matching limits help balance security effectiveness with system performance. When processing high volumes of traffic, limiting the number of pattern states analyzed per packet can prevent performance degradation while maintaining essential intrusion detection capabilities.

Before you begin

Follow these steps to limit pattern matching for intrusions:

Procedure

Step 1 In the access control policy editor, click **Advanced** (**Policies** > **Access Control heading** > **Access Control**, click **Edit** and then click **Advanced Settings**).

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** ( ) next to **Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Pattern Matching Limits** in the **Performance Settings** pop-up window.

Step 4 Enter a value for the maximum number of events to queue in the **Maximum Pattern States to Analyze Per Packet** field.

Step 5 To disable the inspection of packets that will be rebuilt into larger streams of data before and after stream reassembly in Snort 2, check the **Disable Content Checks on Traffic Subject to Future Reassembly** check box. Inspection before and after reassembly requires more processing overhead and may decrease performance.

Important
In Snort 3, the **Disable Content Checks on Traffic Subject to Future Reassembly** check box settings are:

- Checked—Indicates detecting TCP payload before reassembly. It includes inspection of packets before and after stream reassembly. This process requires more processing overhead and may decrease performance.
- Unchecked—Indicates detecting TCP payload after reassembly.

Step 6 Click **OK**.

Step 7 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Regular expression limits overrides for intrusion rules

The default regular expression limits ensure a minimum level of performance. Overriding these limits could increase security, but could also significantly impact performance by permitting packet evaluation against inefficient regular expressions.

Caution Do not override default PCRE limits unless you are an experienced intrusion rule writer with knowledge of the impact of degenerative patterns.

Table 156: Regular expression constraint options

Option	Description
Match Limit State	Specifies whether to override **Match Limit**. You have these options: • select **Default** to use the value configured for **Match Limit** • select **Unlimited** to permit an unlimited number of attempts • select **Custom** to specify either a limit of 1 or greater for **Match Limit**, or to specify 0 to completely disable PCRE match evaluations
Match Limit	Specifies the number of times to attempt to match a pattern defined in a PCRE regular expression.
Match Recursion Limit State	Specifies whether to override **Match Recursion Limit**. You have these options: • select **Default** to use the value configured for **Match Recursion Limit** • select **Unlimited** to permit an unlimited number of recursions • select **Custom** to specify either a limit of 1 or greater for **Match Recursion Limit**, or to specify 0 to completely disable PCRE recursions Note that for **Match Recursion Limit** to be meaningful, it must be smaller than **Match Limit**.
Match Recursion Limit	Specifies the number of recursions when evaluating a PCRE regular expression against the packet payload.

Override regular expression limits for intrusion rules

Override the default regular expression limits to customize performance settings for intrusion rules according to your network requirements.

Regular expression limits control the processing of pattern matching in intrusion rules. Modifying these limits can help optimize performance based on your specific environment and traffic patterns.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** ( ) next to **Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Regular Expression Limits** in the **Performance Settings** pop-up window.

Step 4 You can modify any of the options as described in Regular expression limits overrides for intrusion rules, on page 2867.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Per packet intrusion event generation limits

When the intrusion rules engine evaluates traffic against rules, it places the events generated for a given packet or packet stream in an event queue, then reports the top events in the queue to the user interface. When configuring the intrusion event logging limits, you can specify how many events can be placed in the queue and how many are logged, and select the criteria for determining event order within the queue.

Table 157: Intrusion event logging limits options

Option	Description
Maximum events stored per packet	The maximum number of events that can be stored for a given packet or packet stream.
Maximum events logged per packet	The number of events logged for a given packet or packet stream. This cannot exceed the **Maximum Events Stored Per Packet** value.
Prioritize event logging by	The value used to determine event ordering within the event queue. The highest ordered event is reported through the user interface. You can select from: • `priority`, which orders events in the queue by the event priority. • `content_length`, which orders events by the longest identified content match. When events are ordered by content length, rule events always take precedence over decoder and preprocessor events.

Limiting intrusion events generated per packet

This task allows you to control the number of intrusion events generated per packet by configuring intrusion event logging limits in the access control policy performance settings.

Use this task when you need to manage intrusion detection system performance by limiting the number of events generated for each packet processed by the firewall.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** ( ) next to **Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Intrusion Event Logging Limits** in the **Performance Settings** pop-up window.

Step 4 You can modify any of the options in Per packet intrusion event generation limits, on page 2869.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Packet and intrusion rule latency threshold configuration

Packet and intrusion rule latency threshold configuration is a performance management mechanism that

- uses thresholding to manage packet and rule processing performance in access control policies,
- measures elapsed time for packet and rule processing operations, and
- implements protective actions when configurable time thresholds are exceeded.

Processing mechanisms

This configuration includes two distinct processing mechanisms:

- **Packet latency thresholding**: Measures the total elapsed time taken to process a packet by applicable decoders, preprocessors, and rules, and ceases inspection of the packet if the processing time exceeds a configurable threshold.
- **Rule latency thresholding**: Measures the elapsed time each rule takes to process an individual packet, suspends the violating rule along with a group of related rules for a specified time if the processing time exceeds the rule latency threshold a configurable consecutive number of times, and restores the rules when the suspension expires.

Latency-based performance settings

Latency-based performance settings are a system configuration mechanism that

- takes performance settings from the latest intrusion rule update deployed on the system,
- applies settings based on the security level of the network analysis policy (NAP) associated with the access control policy, and
- evaluates the most secure network analysis policy among custom rules when multiple policies are configured.

Policy evaluation logic

The latency settings that are actually applied depend on the security level of the network analysis policy (NAP) associated with the access control policy. Generally, this is the default NAP policy. However, if custom network analysis rules are configured, and if any of these specify a NAP policy that is more secure than the default NAP policy, then latency settings are based on the most secure NAP policy among the custom rules. If the default NAP policy or any custom rules invoke a custom NAP policy, then the security level used in the evaluation is the system-provided base policy on which each custom NAP policy is based.

This policy evaluation logic applies regardless of whether the effective threshold and network analysis configurations are inherited or configured directly in the policy.

Packet latency thresholding

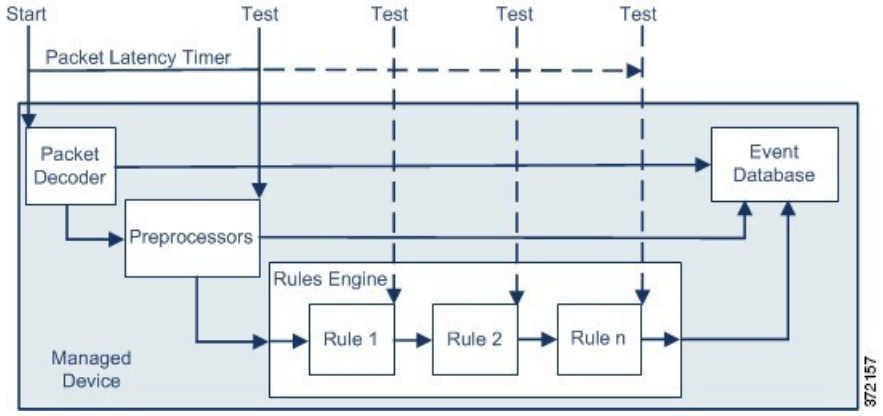
Packet latency thresholding is a software-based latency implementation that

- measures elapsed time, not just processing time, to more accurately reflect the actual time required for the rule to process a packet,
- does not enforce strict timing, and
- can improve system performance by stopping inspection of packets that require excessive processing time.

Packet latency thresholding operation

A timer starts for each packet when decoder processing begins. Timing continues either until all processing ends for the packet or until the processing time exceeds the threshold at a timing test point.

The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks.



Packet latency timing is tested at these test points:

- after the completion of all decoder and preprocessor processing and before rule processing begins
- after processing by each rule

If the processing time exceeds the threshold at any test point, packet inspection ceases.

Tip Total packet processing time does not include routine TCP stream or IP fragment reassembly times.

Packet latency thresholding has no effect on events triggered by a decoder, preprocessor, or rule processing the packet. Any applicable decoder, preprocessor, or rule triggers normally until a packet is fully processed, or until packet processing ends because the latency threshold is exceeded, whichever comes first. If a drop rule detects an intrusion in an inline deployment, the drop rule triggers an event and the packet is dropped.

Note No packets are evaluated against rules after processing for that packet ceases because of a packet latency threshold violation. A rule that would have triggered an event cannot trigger that event, and for drop rules, cannot drop the packet.

Packet latency thresholding can improve system performance in both passive and inline deployments, and can reduce latency in inline deployments, by stopping inspection of packets that require excessive processing time. These performance benefits might occur when, for example:

- for both passive and inline deployments, sequential inspection of a packet by multiple rules requires an excessive amount of time
- for inline deployments, a period of poor network performance, such as when someone downloads an extremely large file, slows packet processing

In a passive deployment, stopping the processing of packets might not contribute to restoring network performance because processing simply moves to the next packet.

Packet latency thresholding notes

By default, the latency-based performance settings for packet handling are disabled. You may choose to enable it. However, Cisco recommends that you do not change the default value for the threshold setting. The information in this section applies only if you choose to specify custom values.

Enable packet latency thresholding

Enable packet latency thresholding to monitor and manage packet processing delays, which helps optimize network performance and identify potential bottlenecks.

Packet latency thresholding allows you to set performance parameters for packet processing. This feature is useful for monitoring network performance and ensuring optimal packet flow through the system.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** ( ) next to **Latency-Based Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.

Step 3 Click **Packet Handling** in the **Latency-Based Performance Settings** pop-up window.

Step 4 Check the **Enabled** check box.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Configure packet latency thresholding

This task enables you to configure packet latency thresholding for latency-based performance settings.

By default, the latency-based performance settings for packet handling is disabled. You may choose to enable it. However, Cisco recommends that you do not change the default value for the threshold setting.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** ( ) next to **Latency-Based Performance Settings**.

System (⚙) > **Monitoring** > **Statistics**

Step 3 If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 Click **Packet Handling** in the **Latency-Based Performance Settings** pop-up window.

By default, **Installed Rule Update** is selected. We recommend using this default.

The values displayed do not reflect the automated settings.

Step 5 If you choose to specify custom values:

- Check the **Enabled** check box, and see Packet latency thresholding notes, on page 2873 for recommended minimum **Threshold** settings.
- You must specify custom values in both the packet handling tab and the rule handling tab.

Step 6 Click **OK**.

Step 7 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Rule latency thresholding

Rule latency thresholding is a software-based latency implementation that

- measures elapsed time, not just processing time, to more accurately reflect the actual time required for the rule to process a packet,
- does not enforce strict timing, and
- temporarily suspends rules that exceed specified processing time thresholds to improve system performance.

Rule latency thresholding operations

A timer measures the processing time each time a packet is processed against a group of rules. Any time the rule processing time exceeds a specified rule latency threshold, the system increments a counter. If the number of consecutive threshold violations reaches a specified number, the system takes these actions:

- suspends the rules for the specified period
- triggers an event indicating the rules have been suspended
- re-enables the rules when the suspension expires
- triggers an event indicating the rules have been re-enabled

The system zeroes the counter when the group of rules has been suspended, or when rule violations are not consecutive. Permitting some consecutive violations before suspending rules lets you ignore occasional rule violations that might have negligible impact on performance and focus instead on the more significant impact of rules that repeatedly exceed the rule latency threshold.

The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks.

Rule latency thresholding has no effect on intrusion events triggered by the rules processing the packet. A rule triggers an event for any intrusion detected in the packet, regardless of whether the rule processing time exceeds the threshold. If the rule detecting the intrusion is a drop rule in an inline deployment, the packet is dropped. When a drop rule detects an intrusion in a packet that results in the rule being suspended, the drop rule triggers an intrusion event, the packet is dropped, and that rule and all related rules are suspended.

Note Packets are not evaluated against suspended rules. A suspended rule that would have triggered an event cannot trigger that event and, for drop rules, cannot drop the packet.

Rule latency thresholding can improve system performance in both passive and inline deployments, and can reduce latency in inline deployments, by suspending rules that take the most time to process packets. Packets are not evaluated again against suspended rules until a configurable time expires, giving the overloaded device time to recover. These performance benefits might occur when:

- hastily written, largely untested rules require an excessive amount of processing time
- a period of poor network performance, such as when someone downloads an extremely large file, causes slow packet inspection

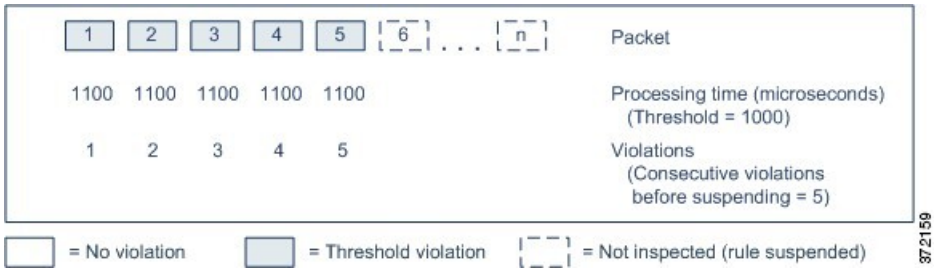
Rule latency thresholding examples

This example shows five consecutive rule processing times that do not result in rule suspension.



In this example, the time required to process each of the first three packets violates the rule latency threshold of 1000 microseconds, and the violations counter increments with each violation. Processing of the fourth packet does not violate the threshold, and the violations counter resets to zero. The fifth packet violates the threshold and the violations counter restarts at one.

This example shows five consecutive rule processing times that do result in rule suspension.



In the second example, the time required to process each of the five packets violates the rule latency threshold of 1000 microseconds. The group of rules is suspended because the rule processing time of 1100 microseconds for each packet violates the threshold of 1000 microseconds for the specified five consecutive violations. Any subsequent packets, represented in the figure as packets 6 through n, are not examined against suspended rules until the suspension expires. If more packets occur after the rules are re-enabled, the violations counter begins again at zero.

Rule latency thresholding notes

Rule latency thresholding suspends rules for the time specified by **Suspension Time** when the time rules take to process a packet exceeds **Threshold** for the consecutive number of times specified by **Consecutive Threshold Violations Before Suspending Rule**.

By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and we recommend that you do not change the default.

The information in this topic applies only if you choose to specify custom values.

You can enable rule 134:1 to generate an event when rules are suspended, and rule 134:2 to generate an event when suspended rules are enabled.

Table 158: Rule latency thresholding options

Option	Description
Threshold	Specifies the time in microseconds that rules should not exceed when examining a packet.
Consecutive Threshold Violations Before Suspending Rule	Specifies the consecutive number of times rules can take longer than the time set for **Threshold** to inspect packets before rules are suspended.
Suspension Time	Specifies the number of seconds to suspend a group of rules.

Configure rule latency thresholding

Configure rule latency thresholding to customize how the system handles rules when latency conditions are detected, improving overall performance and system responsiveness.

By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and we recommend that you do not change the default.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** ( ) next to **Latency-Based Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Rule Handling** in the **Latency-Based Performance Settings** pop-up window.

By default, **Installed Rule Update** is selected. We recommend using this default.

The values displayed do not reflect the automated settings.

Step 4 If you choose to specify custom values:

- You can configure any of the options in Rule latency thresholding notes, on page 2876.
- You must specify custom values in both the packet handling tab and the rule handling tab.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- If you want to generate events, enable latency rules 134:1 and 134:2.
- Deploy configuration changes.

Intrusion performance statistic logging configuration

This reference provides information about intrusion performance statistic logging configuration options, including sample time settings, minimum packet requirements, and troubleshooting options for logging session and protocol distribution data.

Sample time (seconds) and minimum number of packets

When the number of seconds specified elapses between performance statistics updates, the system verifies it has analyzed the specified number of packets. If it has, the system updates performance statistics. Otherwise, the system waits until it analyzes the specified number of packets.

Caution Configuring a very low value, such as one second, for the sample time can greatly impact the device. The device may experience disk space issues, and its operation may be affected by the volume of performance statistics logged. To avoid these issues, do not set a very low value.

Troubleshooting options: log session/protocol distribution

Support might ask you during a troubleshooting call to log protocol distribution, packet length, and port statistics.

Caution Do not enable **Log Session/Protocol Distribution** unless instructed to by Support.

Troubleshooting options: summary

Support might ask you during a troubleshooting call to configure the system to calculate the performance statistics only when the Snort process is shut down or restarted. To enable this option, you must also enable the **Log Session/Protocol Distribution** troubleshooting option.

Caution Do not enable **Summary** unless instructed to do so by Support.

Configure intrusion performance statistic logging

Configure intrusion performance statistic logging to customize how the system samples and logs performance statistics for monitoring and troubleshooting purposes.

Performance statistics logging allows you to monitor device performance by sampling statistics at defined intervals. Proper configuration ensures optimal device performance while providing necessary monitoring data.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to **Performance Settings**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Click **Performance Statistics** in the pop-up window that appears.

Step 3 Modify the **Sample time** or **Minimum number of packets** as described in Intrusion performance statistic logging configuration, on page 2877.

Caution
Configuring a very low value, such as one second, for the **Sample time** can greatly impact the device. The device may experience disk space issues, and its operation may be affected by the volume of performance statistics logged. To avoid these issues, do not set a very low value.

Step 4 Optionally, expand the **Troubleshoot Options** section and modify those options only if asked to do so by Support.

Step 5 Click **OK**.

What to do next

- Deploy configuration changes.



CHAPTER 101

Network Analysis Policies

This chapter provides an insight into network analysis policy basics, prerequisites, and how to manage network analysis policies. It also provides information on custom network analysis policy creation and network analysis policy settings.

- Network analysis policies, on page 2879
- Manage network analysis policies, on page 2880
- Snort 3 definitions and terminologies for network analysis policies, on page 2881
- Custom network analysis policy creation for Snort 3, on page 2883
- Network analysis policy settings and cached changes, on page 2910
- Generate Snort 3 recommendations, on page 2911

Network analysis policies

A *network analysis policy* is a traffic preprocessing configuration that

- governs many traffic preprocessing options and is invoked by advanced settings in your access control policy,
- performs preprocessing after Security Intelligence matching and SSL decryption, and
- completes processing before intrusion or file inspection begins.

Default and custom network analysis policies

By default, the system uses the *Balanced Security and Connectivity* network analysis policy to preprocess all traffic handled by an access control policy. However, you can choose a different default network analysis policy to perform this preprocessing. For your convenience, the system provides a choice of several non-modifiable network analysis policies, which are tuned for a specific balance of security and connectivity by the Talos Intelligence Group. You can also create a custom network analysis policy with custom preprocessing settings.

Note System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules. Network analysis and intrusion policies work together to examine your traffic.

You can also tailor traffic preprocessing options to specific security zones, networks, and VLANs by creating multiple custom network analysis policies, then assigning them to preprocess different traffic.

Manage network analysis policies

This task allows you to perform various management operations on network analysis policies, including creation, editing, deletion, comparison, and report generation.

Under your user name in the toolbar, the system displays a tree of available domains. To switch domains, choose the domain you want to access.

Note Snort 2 is not supported on threat defense version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the management center guide that matches your threat defense version.

Procedure

Step 1 Choose one of these paths to access the network analysis policy.

- **Policies** > **Access Control heading** > **Access Control**, and then click **Network Analysis Policy**
- **Policies** > **Access Control heading** > **Intrusion**, and then click **Network Analysis Policies**

Note
If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Manage your network analysis policy:

- Create—If you want to create a new network analysis policy, click **Create Policy**.

  Two versions of the network analysis policy are created, a **Snort 2 Version** and a **Snort 3 Version**.

  - For the Snort 2 version, see *Custom Network Analysis Policy Creation for Snort 2* in the *Cisco Secure Firewall Management Center Configuration Guide*.
  - For the Snort 3 version, see Custom network analysis policy creation for Snort 3, on page 2883.

- Delete—If you want to delete a network analysis policy, click the **Delete** icon, then confirm that you want to delete the policy. You cannot delete a network analysis policy if an access control policy references it.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Edit—If you want to edit an existing network analysis policy, click the **Edit** icon.

  If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Report—Click the **Report** icon; see *Generating Current Policy Reports* in the *Cisco Secure Firewall Management Center Configuration Guide*.

Snort 3 definitions and terminologies for network analysis policies

This table lists the Snort 3 concepts and terms used in the Network Analysis Policy.

Table 159: Snort 3 definitions and terminologies for network analysis policy

Term	Description
Inspectors	Inspectors are plugins that process packets (similar to the Snort 2 preprocessor).
Binder inspector	The binder inspector defines the flow when a particular inspector must be accessed and considered. When the traffic matches the conditions defined in the binder inspector, only then do the values and configurations for that inspector come into effect. For more information, see *Binder Inspector* in Custom network analysis policy creation for Snort 3, on page 2883.
Singleton inspectors	Singleton inspectors contain one instance. These inspectors do not support adding more instances like multiton inspectors. Settings for a singleton inspector apply to all traffic that matches the inspector, not just a specific traffic segment. For more information, see *Singleton Inspectors* in Custom network analysis policy creation for Snort 3, on page 2883.

Term	Description
Multiton inspectors	Multiton inspectors contain multiple instances which you can configure as needed. These inspectors support configuring settings based on specific conditions, such as network, port, and VLAN. One set of supported settings is called an instance. For more information, see *Multiton Inspectors* in Custom network analysis policy creation for Snort 3, on page 2883.
Schema	The schema file is based on the OpenAPI JSON specification, and it validates the content that you upload or download. You can download the schema file and open it using any third-party JSON editor, such as Swagger editor. The schema file helps you to identify what parameters can be configured for inspectors with their corresponding allowed values, range, and accepted patterns to be used. For more information, see Customize the network analysis policy, on page 2891.
Sample file	It is a pre-existing template that contains example configurations to help you with configuring the inspectors. You can refer to the example configurations included in the sample file and make any changes you require. For more information, see Customize the network analysis policy, on page 2891.
Full configuration	You can download the entire inspector configurations in a single file. All information regarding the inspector configuration is available in this file. The full configuration is a merged configuration of the default configuration (provided as part of the LSP updates by Cisco Talos) and the custom NAP inspector configurations. For more information, see Customize the network analysis policy, on page 2891.

Term	Description
Overridden configuration	In the **Snort 3 Version** of the network analysis policy page: • Under **Actions** > **Upload**, you can click **Overridden Configuration** to upload the JSON file that contains the overridden configuration. • Under **Actions** > **Download**, you can click **Overridden Configuration** to download the inspector configuration that has been overridden. If you have not overridden any inspector configuration, then this option is disabled. When you override the inspector configuration, then this option is enabled automatically to allow you to download. For more information, see Customize the network analysis policy, on page 2891.

Related Topics

Custom network analysis policy creation for Snort 3, on page 2883
Customize the network analysis policy, on page 2891
Network analysis policy mapping, on page 2889

Custom network analysis policy creation for Snort 3

A custom network analysis policy creation for Snort 3 is a data-driven configuration process that

- enables advanced users to customize the default network analysis policy for specific network requirements or to optimize performance, and
- utilizes JSON and JSON Schema based on the OpenAPI specification to configure inspectors and settings.

Network analysis policy customization overview

The default network analysis policy is tuned for typical network requirements and optimal performance. In most situations, you do not need to customize the policy. However, when you have a specific network requirement or when you are facing performance issues, the default network analysis policy can be customized. Note that customizing the network analysis policy is an advanced configuration that should only be done by advanced users or Cisco Support.

In Snort 3, the list of inspectors and settings are not in a one-to-one mapping with the Snort 2 list of preprocessors and settings. The number of inspectors and settings available in Firewall Management Center forms a subset of those that Snort 3 supports. See https://snort.org/snort3 for more information on Snort 3. See https://www.cisco.com/go/snort3-inspectors for more information on the inspectors available in Firewall Management Center.

Note

- While upgrading the Firewall Management Center to the 7.0 release, the changes that were done in the Snort 2 version of the network analysis policy are not migrated to Snort 3 after the upgrade.
- Unlike the intrusion policy, there is no option to synchronize Snort 2 network analysis policy settings to Snort 3.

Default inspector updates

Lightweight Security Package (LSP) updates may contain new inspectors or modifications to integer ranges for existing inspector configurations. After you install an LSP, new inspectors or updated ranges will be available under **Inspectors** in the **Snort 3 Version** of your network analysis policy.

Binder inspector

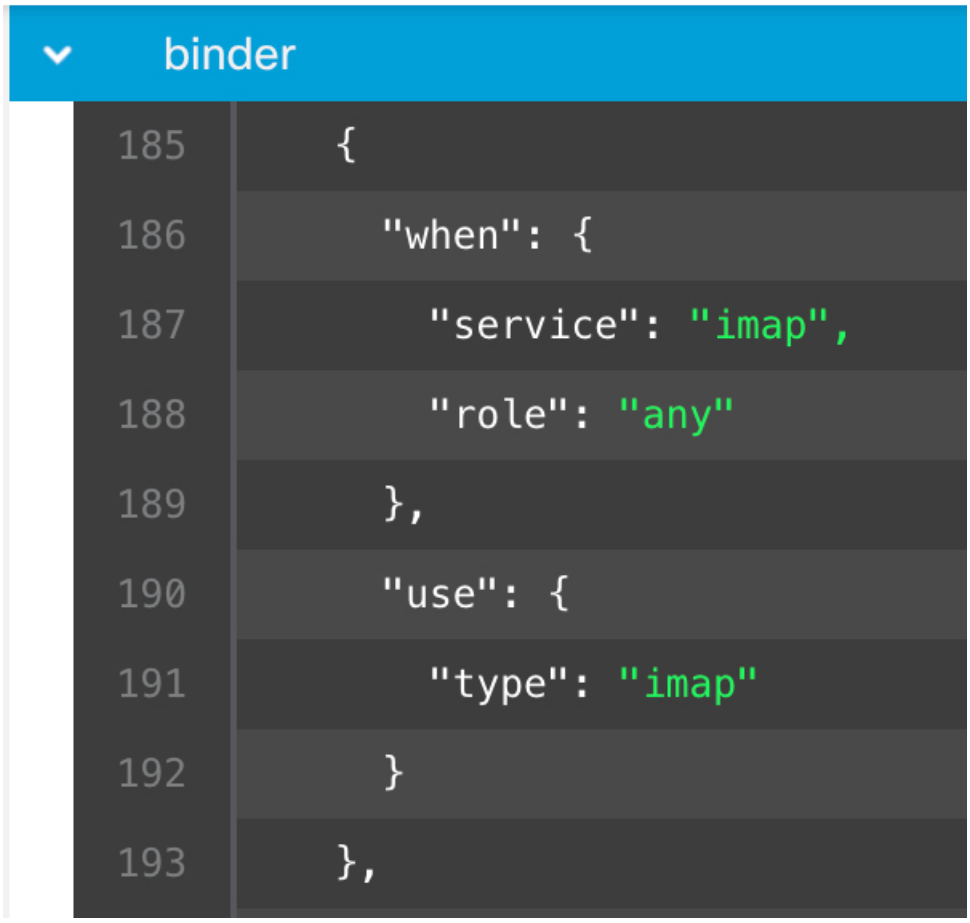
Binder inspector defines the flow when a particular inspector has to be accessed and taken into consideration. When the traffic matches the conditions defined in the binder inspector, only then the values and configurations for that inspector come into effect. For example:

For the *imap* inspector, the binder defines the following condition when it has to be accessed. That is when:

- Service is equal to imap.
- Role is equal to any.

a

If these conditions are met, then use the type imap.



Singleton inspectors

Singleton inspectors contain a single instance. These inspectors do not support adding more instances like multiton inspectors. Settings of a singleton inspector apply to all traffic, not to a specific segment.

For example:

```
{
   "normalizer":{
      "enabled":true,
      "type":"singleton",
      "data":{
         "ip4":{
            "df":true
         }
      }
   }
}
```

Multiton inspectors

Multiton inspectors contain multiple instances which you can configure as needed. These inspectors support configuring settings based on specific conditions, such as network, port, and VLAN. One set of supported settings is called an instance. There is a default instance, and you can also add additional instances based on specific conditions. If the traffic matches that condition, the settings from that instance are applied. Otherwise,

the settings from the default instance are applied. Also, the name of the default instance is the same as the name of the inspector.

For a multiton inspector, when you upload the overridden inspector configuration, you also need to include and define a matching binder condition (conditions under when the inspector has to be accessed or used) for each instance in the JSON file, otherwise, the upload will result in an error. You can also create new instances. Ensure that you include a binder condition for each new instance you create to avoid errors.

For example:

- Multiton inspector where the default instance is modified.

```
{
   "http_inspect":{
      "enabled":true,
      "type":"multiton",
      "instances":[
         {
            "name":"http_inspect",
            "data":{
               "response_depth":5000
            }
         }
      ]
   }
}
```

- Multiton inspector where the default instance and default binder is modified.

```
{
   "http_inspect":{
      "enabled":true,
      "type":"multiton",
      "instances":[
         {
            "name":"http_inspect",
            "data":{
               "response_depth":5000
            }
         }
      ]
   },
   "binder":{
      "type":"binder",
      "enabled":true,
      "rules":[
         {
            "use":{
               "type":"http_inspect"
            },
            "when":{
               "role":"any",
               "ports":"8080",
               "proto":"tcp",
               "service":"http"
            }
         }
      ]
   }
}
```

- Multiton inspector where a custom instance and a custom binder is added.

```
{
   "http_inspect":{
      "enabled":true,
      "type":"multiton",
      "instances":[
         {
            "name":"http_inspect1",
            "data":{
               "response_depth":5000
            }
         }
      ]
   },
   "binder":{
      "type":"binder",
      "enabled":true,
      "rules":[
         {
            "use":{
               "type":"http_inspect",
               "name":"http_inspect1"
            },
            "when":{
               "role":"any",
               "ports":"8080",
               "proto":"tcp",
               "service":"http"
            }
         }
      ]
   }
}
```

Common Industrial Protocol safety

Common Industrial Protocol (CIP) safety is a set of extensions to the CIP that

- enables the safe operation of devices,
- provides fail-safe communication between different nodes on a CIP network, and
- comprises two main components for safety parameter exchange and safety information transmission.

CIP safety protocol components

The CIP safety protocol comprises two main components:

- CIP safety segments—Used in Forward Open messages to exchange safety parameters for the subsequent safety session.
- CIP safety messages—Used to exchange actual safety information.

The CIP inspector detects and identifies:

- CIP as a service and client
- Payloads, such as CIP Read, CIP Admin, CIP Infrastructure, and CIP Write

The CIP inspector can parse the CIP segments and detect the CIP Safety segments in the Forward Open requests.

To test the CIP safety feature, you must enable the CIP inspector. See Detect and block safety segments in CIP packets, on page 2888.

Detect and block safety segments in CIP packets

This task enables detection and blocking of CIP safety segments while allowing other CIP packets to pass through.

Use this procedure to detect and block CIP safety segments while allowing other CIP packets:

- Create a custom network analysis policy called **cip_safety**.
- Create access control rules in your access control policy to block CIP safety and to allow all other packets.

To test the CIP safety feature, enable the CIP inspector in the management center and assign it to an access control policy.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Click the **Snort 3 Version** of the network analysis policy **cip_safety** that you created.

Step 3 Under **Inspectors**, click **CIP** to expand it.

The default configuration appears in the left column and the overridden configuration appears in the right column under the inspector.

Step 4 Under **Overridden Configuration** on the right column, click the **Edit Inspector** icon and change the "enabled" field in **CIP** from false (default) to true.

Step 5 Click **OK** and then click **Save**.

Step 6 To assign the **CIP** inspector to the access control policy, choose **Policies** > **Access Control heading** > **Access Control**, click **Edit** and choose the **Advanced Settings** option from the **More** drop-down arrow at the end of the packet flow line.

Step 7 Click **Edit** () next to **Network Analysis and Intrusion Policies**.

Step 8 In the **Network Analysis and Intrusion Policies** window, choose the access control policy **cip_safety** that you created from the **Default Network Analysis Policy** drop-down list.

The CIP inspector is now enabled in the management center and you can create the custom access control rules to block CIP safety and to allow all other CIP packets.

Step 9 After you send live traffic containing CIP safety packet flows, go to **Unified Events** to verify that the payload is the expected payload that contains CIP safety packet logs for the detection and block use case as mentioned in this procedure.

CIP is detected as an application protocol and client (see the **Application Protocol** and **Client Application** fields), and **CIP Safety** is shown under the **Web Application** field.

Network analysis policy mapping

Network analysis policy mapping is a correlation mechanism that

- provides mapping information for network analysis policies,
- finds the corresponding Snort 2 version of policies for the Snort 3 version, and
- ensures that the Snort 3 version of policies has its equivalent Snort 2 version.

View network analysis policy mapping

View which Snort 3 network analysis policies are automatically mapped to their Snort 2 equivalent policies in the system.

Network analysis policy mapping allows you to see the relationship between Snort 3 and Snort 2 policies when transitioning between versions or managing both simultaneously.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Click **NAP Mapping**.

Step 3 Expand the arrow for **View Mappings**.

The Snort 3 network analysis policies that are automatically mapped to a Snort 2 equivalent policy are displayed.

Step 4 Click **OK**.

Create a network analysis policy

Create a new network analysis policy that provides both Snort 2 and Snort 3 versions for network traffic inspection and analysis.

All the existing network analysis policies are available in Firewall Management Center with their corresponding Snort 2 and Snort 3 versions. When you create a new network analysis policy, it is created with both the Snort 2 version and the Snort 3 version.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Click **Create Policy**.

Step 3 Enter the **Name** and **Description**.

Step 4 Select a **Base Policy** and click **Save**.

The new network analysis policy is created with its corresponding **Snort 2 Version** and **Snort 3 Version**.

Modify the network analysis policy

This task allows you to update the configuration of an existing network analysis policy by modifying its properties such as name, description, or base policy settings.

You can modify the network analysis policy to change its name, description, or the base policy.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Click **Edit** to change the name, description, inspection mode, or the base policy.

Note
If you edit the network analysis policy name, description, base policy, and inspection mode, the edits are applied to both the Snort 2 and Snort 3 versions. If you want to change the inspection mode for a specific version, then you can do that from within the network analysis policy page for that respective version.

Step 3 Click **Save**.

Search for an inspector on the Network Analysis Policy window

This task allows you to quickly locate specific inspectors within the network analysis policy interface by using search functionality.

On the Snort 3 version of the network analysis policy window, you may need to search for an inspector by entering any relevant text in the search bar.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Enter the name of an inspector or any relevant text to search for in the **Search** bar.

All the inspectors matching the text you search for are displayed.

For example, if you enter `pop`, then the pop inspector and the binder inspector are shown as matching results on the screen.

Copy the inspector configuration

This task describes how to copy inspector configurations.

You can copy the inspector configuration for the Snort 3 version of the network analysis policy according to your requirements.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to copy the configuration.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

Step 4 Click the **Copy to clipboard** icon to copy the inspector configuration to the clipboard for one or both of these options.

- **Default Configuration** in the left column
- **Overridden Configuration** in the right column

Step 5 Paste the copied inspector configuration to a JSON editor to make any edits you may require.

Related Topics

Customize the network analysis policy, on page 2891

Customize the network analysis policy

Customize the network analysis policy to adapt Snort 3 configurations according to your security requirements.

You can customize the Snort 3 version of the network analysis policy according to your requirements.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

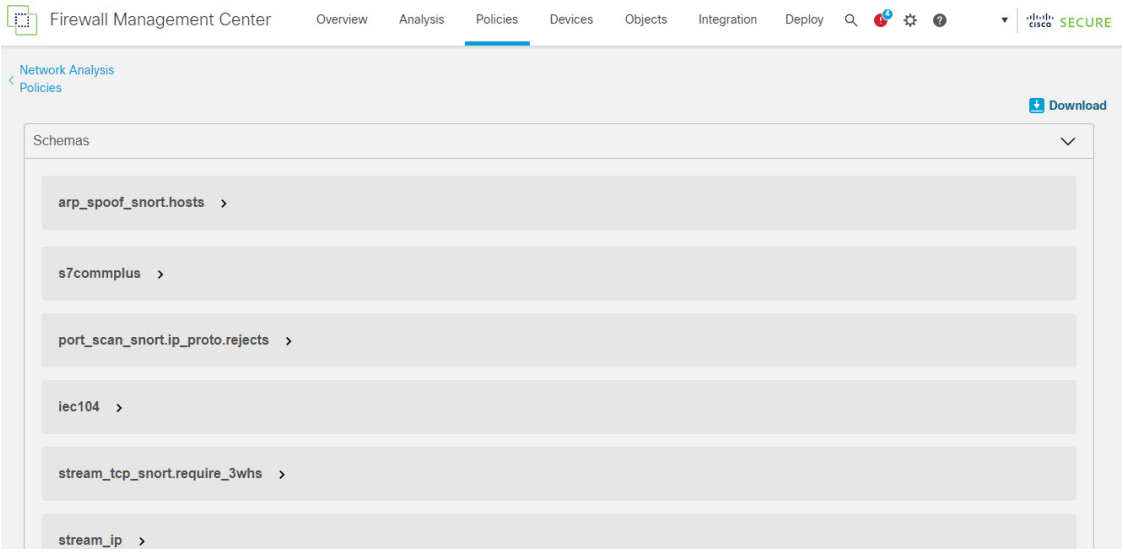
Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Click the **Actions** dropdown menu.

The following options are displayed:

- View Schema
- Download Schema / Download Sample File / Template
- Download Full Configuration
- Download Overridden Configuration
- Upload Overridden Configuration

Step 4 Click **View Schema** to open the schema file directly in a browser.



Step 5 You can download the schema file, sample file / template, full configuration, or overridden configuration as needed.

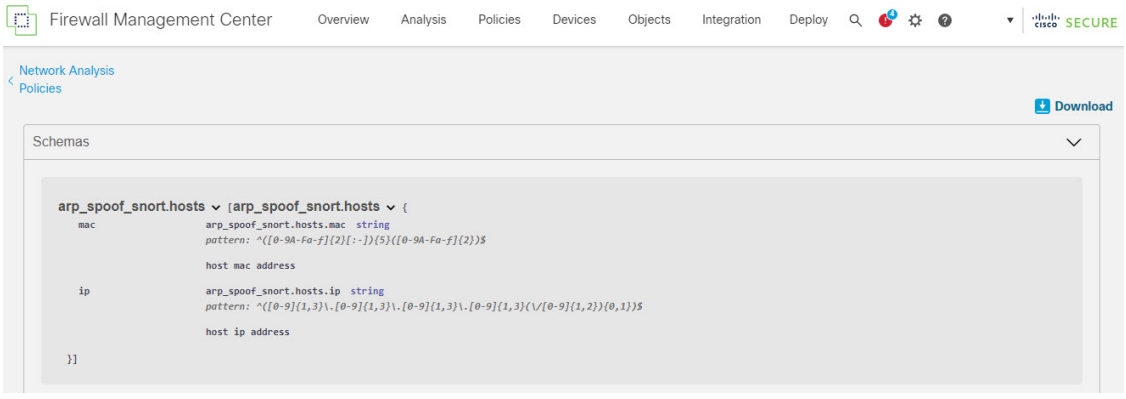
These options provide you an insight about the allowed values, range, and patterns, existing and default inspector configurations, and overridden inspector configurations.

a) Click **Download Schema** to download the schema file.

The schema file validates the content that you upload or download. You can download the schema file and open it using any third-party JSON editor. The schema file helps you identify which parameters can be configured for inspectors, along with their allowed values, ranges, and accepted patterns.

For example, for the *arp_spoof_snort* inspector, you can configure the hosts. The hosts include the *mac* and *ip* address values. The schema file shows the following accepted pattern for these values.

- **mac** – `pattern: ^([0-9A-Fa-f]{2}[:-]){5}([0-9A-Fa-f]{2})$`
- **ip** – `pattern: ^([0-9]{1,3}.[0-9]{1,3}.[0-9]{1,3}.[0-9]{1,3}(/[0-9]{1,2}){0,1})$`



You must provide the values, range, patterns according to the accepted ones in the schema file to be able to successfully override the inspector configuration, otherwise, you get an error message.

b) Click **Download Sample File / Template** to use a pre-existing template that contains example configurations to help you with configuring the inspectors.

You can refer to the example configurations included in the sample file and make any changes that you may require.

c) Click **Download Full Configuration** to download the entire inspector configurations in a single JSON file.

You can download the full configuration to find the information you need, rather than expanding the inspectors separately. This file contains all information about the inspector configuration.

d) Click **Download Overridden Configuration** to download the inspector configuration that has been overridden.

Step 6 To override the existing configuration, follow these steps.

You can choose to override an inspector configuration using these ways.

- Make inline edits for an inspector directly on the Firewall Management Center. See the topic **Make Inline Edit for an Inspector to Override Configuration** in the **Getting Started with Network Analysis Policies** chapter of the *Cisco Secure Firewall Management Center Snort 3 Configuration Guide*.
- Continue to follow the current procedure of using the **Actions** dropdown menu to upload the overridden configuration file.

If you chose to make inline edits directly in the Firewall Management Center, then you do not need to complete this procedure any further. Otherwise, you must complete this procedure completely.

a) Under **Inspectors**, expand the required inspector for which you want to override the default configuration.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

You may need to search for an inspector by entering any relevant text in in the search bar.

b) Click the **Copy to clipboard** icon to copy the default inspector configuration to the clipboard.
c) Create a JSON file and paste the default configuration in it.
d) Keep the inspector configuration that you want to override, and remove all other configuration and instances from the JSON file.

You can also use the **Sample File / Template** to understand how to override the default configuration. This is a sample file that includes JSON snippets explaining how you can customize the network analysis policy for Snort 3.

e) Make changes to the inspector configuration as needed.

Validate the changes and make sure they conform to the schema file. For multiton inspectors, make sure that the binder conditions for all instances are included in the JSON file. See *Multiton Inspectors* in the topic **Custom Network Analysis Policy Creation for Snort 3** in the *Cisco Secure Firewall Management Center Snort 3 Configuration Guide* for more information.

f) If you are copying any further default inspector configurations, add the inspector configuration to the existing file that contains the overridden configuration.

Note
The copied inspector configuration must comply with the JSON standards.

g) Save the overridden configuration file to your system.

Step 7 From the **Actions** dropdown menu, choose Upload Overridden Configuration to upload the JSON file that contains the overridden configuration.

Caution
Upload only the changes that you require. You should not upload the entire configuration as it makes the overrides sticky in nature and therefore, any subsequent changes to the default configuration as part of the LSP updates would not be applied.

You can drag and drop a file or click to browse to the JSON file saved in your system that contains the overridden inspector configuration.

- **Merge inspector overrides** – Content in the uploaded file is merged with the existing configuration if there is no common inspector. If there are common inspectors, then the content in the uploaded file (for common inspectors) takes precedence over the previous content, and it replaces the previous configuration for those inspectors.
- **Replace inspector overrides** – Removes all previous overrides and replaces them with the new content in the uploaded file.

  Attention
  Choosing this option deletes all the previous overrides. Review your choice carefully before using this option to override the configuration.

If any error occurs while uploading the overridden inspectors, you see the error in the **Upload Overridden Configuration File** pop-up window. You can also download the file with the error, fix the error, and reupload the file.

Step 8 In the **Upload Overridden Configuration File** pop-up window, click **Import** to upload the overridden inspector configuration.

After you upload the overridden inspector configuration, you will see an orange icon next to the inspector that signifies that it is an overridden inspector.

Also, the **Overridden Configuration** column under the inspector shows the overridden value.

You can also view all the overridden inspectors using the **Show Overrides Only** checkbox adjacent to the Search bar.

Note
Make sure that you always download the overridden configuration, open the JSON file, and append any new changes/overrides to the inspector configurations to this file. This action is needed so that you do not lose the old overridden configurations.

Step 9 (Optional) Take a backup of the overridden configuration file on your system before making any new inspector configuration changes.

Tip
We recommend that you take a backup from time to time as you override the inspector configuration.

Make inline edit for an inspector to override configuration

This task allows you to make inline edits for inspector configurations to override the default configuration according to your requirements in the Snort 3 version of the network analysis policy.

For the Snort 3 version of the network analysis policy, you can make an inline edit for the inspector configuration to override the configuration according to your requirements.

Alternatively, you can also use the **Actions** drop-down menu to upload the overridden configuration file. See Customize the network analysis policy, on page 2891 for more information.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to override the default setting.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

Step 4 Under the **Overridden Configuration** in the right column, click **Edit Inspector** (Pencil) icon to make changes to the inspector configuration.

The Override Configuration pop-up appears where you can make the required edits.

Note

- Make sure that you keep only those settings that you want to override. If you leave a setting with the same value, that field becomes sticky. This means if that setting is changed in the future by Talos, the current value will be retained.
- If you are adding or deleting any custom instance, make sure that you add or delete a binder rule for that instance in the binder inspector as well.

Step 5 Click **OK**.

If there are any errors according to the JSON standards, it shows you an error message.

Step 6 Click **Save** to save the changes.

If the changes conform to the OpenAPI schema specification, the Firewall Management Center allows you to save the configuration, otherwise, the **Error saving overridden configuration** pop-up appears that shows the errors. You can also download the file with the errors.

Related Topics

Customize the network analysis policy, on page 2891
Revert unsaved changes during inline edits, on page 2895
Revert overridden configuration to default configuration, on page 2897
Examples of custom network analysis policy configuration, on page 2899

Revert unsaved changes during inline edits

This task allows you to revert any unsaved changes made during inline edits to override the configuration for an inspector, returning to the most recently saved value.

While making inline edits to override the configuration for an inspector , you can revert any unsaved changes. Note that this action reverts all unsaved changes to the most recently saved value, but does not revert the configuration to the default configuration for an inspector.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to revert the unsaved changes.

The default configuration is displayed in the left column and the overridden configuration is displayed in the right column under the inspector.

Step 4 Under the **Overridden Configuration** on the right column, click the **Cross** ✕ icon to revert any unsaved changes for the inspector.

Alternatively, you can click **Cancel** to cancel the changes.

If you do not have any unsaved changes to the inspector configuration, then this option is not visible.

Related Topics

Revert overridden configuration to default configuration, on page 2897
Make inline edit for an inspector to override configuration, on page 2894

View the list of inspectors with overrides

This task enables you to identify and review all inspectors that have been overridden in your network analysis policy configuration.

You can view a list of all the overridden inspectors.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Check the **Show Overrides Only** checkbox adjacent to the Search bar to view the list of overridden inspectors.

All the overridden inspectors are shown with an orange icon next to their names to help you identify them.

Related Topics

Search for an inspector on the Network Analysis Policy window, on page 2890
Make inline edit for an inspector to override configuration, on page 2894
Customize the network analysis policy, on page 2891

Revert overridden configuration to default configuration

Revert any changes that you made to override the default configuration for an inspector, restoring the original default settings.

You can revert any changes that you made to override the default configuration for an inspector. This action reverts the overridden configuration to the default configuration for an inspector.

Procedure

Step 1 Go to **Policies** > **Access Control heading** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to revert the overridden configuration.

The overridden inspectors are shown with the orange icon next to their name.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector. Under the **Overridden Configuration** on the right column, click **Revert to default configuration** (back arrow) icon to revert the overridden configuration for the inspector to the default configuration.

If you did not make any changes to the default configuration for the inspector, then this option is disabled.

Step 4 Click **Revert** to confirm the decision.

Step 5 Click **Save** to save the changes.

If you do not want to save the changes, you can click **Cancel** or the **Cross** ✕ icon.

Related Topics

Revert unsaved changes during inline edits, on page 2895
Customize the network analysis policy, on page 2891
Make inline edit for an inspector to override configuration, on page 2894
Examples of custom network analysis policy configuration, on page 2899

Validate Snort 3 policies

This task validates Snort 3 policies to ensure they work correctly with current and previous device versions, handling version-specific configurations and inspector settings.

To validate the Snort 3 policies, here is a list of basic information that you can make note of:

- Current version of the Cloud-Delivered Firewall Management Center can manage multiple Firewall Threat Defense versions.
- Current version of Cloud-Delivered Firewall Management Center supports network analysis policy (NAP) configurations which are not applicable to previous version of Firewall Threat Defense devices.
- Current NAP Policy and validations will work based on the current version support.
- Changes may include content which is not valid for previous versions of Firewall Threat Defenses.

- Policy configuration changes are accepted if they are valid for the current version and performed using the current Snort 3 binary and NAP schema.
- For previous version Firewall Threat Defenses, validation is performed during deployment using NAP schema and Snort 3 binary for that specific version. If there is any configuration which is not applicable for the given version, the user is informed or warned that the configuration not supported on the given version will not be deployed, while the remaining configuration will be deployed.

In this procedure, when we associate the NAP policy to an Access Control Policy and deploy it on a device, for example any inspector like rate filter configuration is applied to validate the Snort 3 policies.

Procedure

Step 1 **Steps to Override NAP Policy Configuration:** Under **Inspectors** in the **Snort 3 Version** of the network analysis policy, expand the required inspector for which you want to override the default setting.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

a) Under the **Overridden Configuration** on the right column, click **Edit Inspector** (Pencil) icon to make changes to any inspector like rate_filter.

The Override Configuration pop-up appears where you can make the required edits to the rate_filter inspector. Click **OK**.

b) Click **Save** to save the changes.

Alternatively, you can use the **Actions** dropdown menu to upload the overridden configuration file.

Step 2 Click the **Actions** dropdown menu in the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Upload**, you can click **Overridden Configuration** to upload the JSON file that contains the overridden configuration.

Caution
Upload only the changes that you require. You should not upload the entire configuration as it causes the overrides to remain persistent, so any subsequent changes to the default configuration from LSP updates will not be applied.

You can drag and drop a file or click to browse to the JSON file saved in your system that contains the overridden inspector configuration.

- **Merge inspector overrides**: Content in the uploaded file is merged with the existing configuration if there is no common inspector. If there are common inspectors, then the content in the uploaded file (for common inspectors) takes precedence over the previous content, and it replaces the previous configuration for those inspectors.
- **Replace inspector overrides**: Removes all previous overrides and replaces them with the new content in the uploaded file.

Note
As choosing this option deletes all the previous overrides, make an informed decision before overriding the configuration using this option.

If any error occurs while uploading the overridden inspectors, you see the error on the **Upload Overridden Configuration File** pop-up window. You can also download the file with the error, then fix the error and reupload the file.

Step 4 **Steps to Associate NAP Policy to Access Control Policy:** In the access control policy editor, click **Advanced**, then click **Edit** next to the Network Analysis and Intrusion Policies section.

a) From the **Default Network Analysis Policy** drop-down list, select a default network analysis policy.

If you choose a user-created policy, you can click **Edit** to edit the policy in a new window. You cannot edit system-provided policies.

b) Click **OK** and then click **Save** to save the policy.

Step 5 Alternatively, in the access control policy editor, click **Advanced**, then click **Edit** next to the Network Analysis and Intrusion Policies section.

a) Click **Add Rule**.
b) Configure the rule's conditions by clicking the options you want to add.
c) Click **Network Analysis** and choose the **Network Analysis Policy** you want to use to preprocess the traffic matching this rule.
d) Click **Add**.

Step 6 **Deployment:** On the Cloud-Delivered Firewall Management Center menu bar, click **Deploy** and then select **Deployment**.

Step 7 Identify and choose the devices on which you want to deploy configuration changes.

- Search: Search for the device name, type, domain, group, or status in the search box.
- Expand: Click **Expand Arrow** to view device-specific configuration changes to be deployed.

By selecting the device check box, all changes for that device listed under it are pushed for deployment. Alternatively, you can use the **Policy Selection** to select individual policies or configurations to deploy, withholding the rest.

Optionally, use **Show or Hide Policy** to selectively view or hide the associated unmodified policies.

Step 8 Click **Deploy**.

Step 9 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

Note
It displays a warning that Snort 3 network analysis policy contains inspectors or attributes not valid for this Firewall Threat Defense version. Invalid settings will be skipped during deployment. Invalid inspectors are: ["rate_filter"] only for devices lower than 7.1 version.

Examples of custom network analysis policy configuration

This is a sample file that includes JSON snippets explaining how you can customize the network analysis policy for Snort 3. You can choose to override an inspector configuration using the following ways:

- Make inline edits for an inspector directly on the Firewall Management Center. See Make inline edit for an inspector to override configuration, on page 2894.

- Use the **Actions** drop-down menu to upload the overridden configuration file. See Customize the network analysis policy, on page 2891.

Before you choose any of these options, review all the details and examples to help you define the network analysis policy overrides successfully. To avoid risks and errors, read and understand the scenarios provided here.

If you choose to override an inspector configuration from the **Actions** drop-down menu, you need to construct a JSON file for the network analysis policy overrides and upload the file.

For overriding an inspector configuration in the network analysis policy, you must upload only the changes that you require. You should not upload the entire configuration because it makes the overrides sticky in nature and therefore, any subsequent changes to the default values or configuration as part of the LSP updates would not be applied.

Here are the examples for various scenarios:

Enabling a Singleton Inspector when the Default State in the Base Policy is Disabled

```
{
  "rate_filter": {
    "enabled": true,
    "type": "singleton",
    "data": []
  }
}
```

Disabling a Singleton Inspector when the Default State in the Base Policy is Enabled

```
{
  "rate_filter": {
    "enabled": false,
    "type": "singleton",
    "data": []
  }
}
```

Enabling a Multiton Inspector when the Default State in the Base Policy is Disabled

```
{
  "ssh": {
    "enabled": true,
    "type": "multiton",
    "instances": []
  }
}
```

Disabling a Multiton Inspector when the Default State in the Base Policy is Enabled

```
{
  "ssh": {
    "enabled": false,
    "type": "multiton",
    "instances": []
  },
  "iec104": {
    "type": "multiton",
    "enabled": false,
    "instances": []
```

```
  }
}
```

Overriding the Default Value of Specific Setting(s) for Singleton Inspector

```
{
  "normalizer": {
    "enabled": true,
    "type": "singleton",
    "data": {
      "tcp": {
        "block": true
      },
      "ip6": true
    }
  }
}
```

Overriding Specific Setting(s) of a Default Instance (where Instance Name Matches with Inspector Type) in Multiton Inspector

```
{
  "http_inspect": {
    "enabled": true,
    "type": "multiton",
    "instances": [
      {
        "data": {
          "unzip": false
        },
        "name": "http_inspect"
      }
    ]
  }
}
```

Adding Binder Rule for a Default Instance with Required Changes



Note Default binder rules cannot be edited; they are always appended at the end.

```
{
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "use": {
          "type": "http_inspect"
        },
        "when": {
          "role": "server",
          "service": "http",
          "dst_nets": "10.1.1.0/24"
        }
      }
    ]
  }
}
```

Adding a New Custom Instance



Note Corresponding binder rule entry must be defined in the binder inspector.

```
{
  "telnet": {
    "enabled": true,
    "type": "multiton",
    "instances": [
      {
        "name": "telnet_my_instance",
        "data": {
          "encrypted_traffic": true
        }
      }
    ]
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_my_instance"
        }
      }
    ]
  }
}
```

Overriding a Singleton Instance, Multiton Default Instance, and Creating a New Multiton Instance in a Single JSON Override

Example to show the following in a single JSON override:

- Overriding a Singleton instance (**`normalizer`** inspector)
- Overriding a Multiton default instance (**`http_inspect`** inspector)
- Creating a new Multiton instance (**`telnet`** inspector)

```
{
  "normalizer": {
    "enabled": true,
    "type": "singleton",
    "data": {
      "tcp": {
        "block": true
      },
      "ip6": true
    }
  },
  "http_inspect": {
    "enabled": true,
```

```
      "type": "multiton",
      "instances": [
        {
          "data": {
            "unzip": false,
            "xff_headers": "x-forwarded-for true-client-ip x-another-forwarding-header"
          },
          "name": "http_inspect"
        }
      ]
    },
    "telnet": {
      "enabled": true,
      "type": "multiton",
      "instances": [
        {
          "name": "telnet_my_instance",
          "data": {
            "encrypted_traffic": true
          }
        }
      ]
    },
    "binder": {
      "enabled": true,
      "type": "binder",
      "rules": [
        {
          "when": {
            "role": "any",
            "service": "telnet"
          },
          "use": {
            "type": "telnet",
            "name": "telnet_my_instance"
          }
        },
        {
          "use": {
            "type": "http_inspect"
          },
          "when": {
            "role": "server",
            "service": "http",
            "dst_nets": "10.1.1.0/24"
          }
        }
      ]
    }
}
```

Note You do not need to provide the **name** attribute for the default instance in binder rules.

Configuring arp_spoof

Example for configuring **arp_spoof**:

The **arp_spoof** inspector does not have any default configurations for any attributes. This demonstrates the case where you can provide the overrides.

```
{
  "arp_spoof": {
    "type": "singleton",
    "data": {
      "hosts": [
        {
          "ip": "1.1.1.1",
          "mac": "ff:0f:f1:0f:0f:ff"
        },
        {
          "ip": "2.2.2.2",
          "mac": "ff:0f:f2:0f:0f:ff"
        }
      ]
    },
    "enabled": true
  }
}
```

Configuring rate_filter

```
{
  "rate_filter": {
    "data": [
      {
        "apply_to": "[10.1.2.100, 10.1.2.101]",
        "count": 5,
        "gid": 135,
        "new_action": "alert",
        "seconds": 1,
        "sid": 1,
        "timeout": 5,
        "track": "by_src"
      }
    ],
    "enabled": true,
    "type": "singleton"
  }
}
```

Configuring Binder Rules when Multi-Hierarchy Network Analysis Policy is Used

This example illustrates adding a new custom instance in child policy and the way binder rules should be written. Binder rules are defined as a list and therefore, it is important to pick up the rules defined in the parent policy and build the new rules on top of it as rules will not be merged automatically. The binder rules available in child policy are a source of truth in totality.

On the Firewall Threat Defense, the default Cisco Talos policy rules are appended on these user-defined overrides.

Parent Policy:

We have defined a custom instance by the name **`telnet_parent_instance`** and the corresponding binder rule.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": true
```

```
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

Child Policy:

This network analysis policy has the aforementioned policy as its base policy. We have defined a custom instance by the name **telnet_child_instance** and have also defined the binder rules for this instance. The binder rules from parent policy need to be copied here, and then child policy binder rules can be prepended or appended on top of it based on the nature of the rule.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": false
        },
        "name": "telnet_child_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet",
          "nets": "10.2.2.0/24"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_child_instance"
        }
      },
      {
        "when": {
          "role": "any",
          "service": "telnet"
```

```
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

Configuring List Inspector Attribute in General

While changing overrides for any attribute of type list, it is important to pass the full contents rather than partial override. This means if a base policy attributes are defined as:

```
{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1"
      }
    },
    {
      "entry2": {
        "key2": "value2"
      }
    }
  ]
}
```

If you want to modify **value1** to **value1-new**, the override payload must look like the following:

Correct Way:

```
{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1-new"
      }
    },
    {
      "entry2": {
        "key2": "value2"
      }
    }
  ]
}
```

Incorrect Way:

```
{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1-new"
      }
    }
  ]
}
```

You can understand this configuration by taking the trimmed values of the **alt_max_command_line_len** attribute in the **smtp** inspector. Suppose the default (base) policy configuration for **smtp** inspector is as follows:

```
{
  "smtp": {
    "type": "multiton",
    "instances": [
      {
        "name": "smtp",
        "data": {
          "decompress_zip": false,
          "normalize_cmds": "ATRN AUTH BDAT CHUNKING DATA DEBUG EHLO
           EMAL ESAM ESND ESOM ETRN EVFY EXPN HELO HELP IDENT MAIL
           NOOP ONEX QUEU QUIT RCPT RSET SAML SEND SOML STARTTLS TICK
           TIME TURN TURNME VERB VRFY X-ADAT XADR XAUTH XCIR X-DRCP X-
           ERCP XEXCH50 X-EXCH50 X-EXPS XGEN XLICENSE X-LINK2STATE XQUE
           XSTA XTRN XUSR",
          "ignore_data": false,
          "max_command_line_len": 512,
          "max_header_line_len": 1000,
          "log_rcptto": false,
          "decompress_swf": false,
          "max_response_line_len": 512,
          "b64_decode_depth": -1,
          "max_auth_command_line_len": 1000,
          "log_email_hdrs": false,
          "xlink2state": "alert",
          "binary_data_cmds": "BDAT XEXCH50",
          "auth_cmds": "AUTH XAUTH X-EXPS",
          "log_filename": false,
          "uu_decode_depth": -1,
          "ignore_tls_data": false,
          "data_cmds": "DATA",
          "bitenc_decode_depth": -1,
          "alt_max_command_line_len": [
            {
              "length": 255,
              "command": "ATRN"
            },
            {
              "command": "AUTH",
              "length": 246
            },
            {
              "length": 255,
              "command": "BDAT"
            },
            {
              "length": 246,
              "command": "DATA"
            }
          ],
          "log_mailfrom": false,
          "decompress_pdf": false,
          "normalize": "none",
          "email_hdrs_log_depth": 1464,
          "valid_cmds": "ATRN AUTH BDAT CHUNKING DATA DEBUG EHLO
           EMAL ESAM ESND ESOM ETRN EVFY EXPN HELO HELP IDENT MAIL
           NOOP ONEX QUEU QUIT RCPT RSET SAML SEND SOML STARTTLS TICK
           TIME TURN TURNME VERB VRFY X-ADAT XADR XAUTH XCIR X-DRCP X-
           ERCP XEXCH50 X-EXCH50 X-EXPS XGEN XLICENSE X-LINK2STATE XQUE
           XSTA XTRN XUSR",
          "qp_decode_depth": -1
        }
      }
    ],
    "enabled": true
```

```
  }
}
```

Now, if you want to add two more objects to the **alt_max_command_line_len** list:

```
{
    "length": 246,
    "command": "XEXCH50"
},
{
    "length": 246,
    "command": "X-EXPS"
}
```

Then the custom network analysis policy override JSON would look like the following:

```
{
  "smtp": {
    "type": "multiton",
    "instances": [
      {
        "name": "smtp",
        "data": {
          "alt_max_command_line_len": [
            {
              "length": 255,
              "command": "ATRN"
            },
            {
              "command": "AUTH",
              "length": 246
            },
            {
              "length": 255,
              "command": "BDAT"
            },
            {
              "length": 246,
              "command": "DATA"
            },
            {
              "length": 246,
              "command": "XEXCH50"
            },
            {
              "length": 246,
              "command": "X-EXPS"
            }
          ]
        }
      }
    ],
    "enabled": true
  }
}
```

Configuring Overrides when Multi-Hierarchy Network Analysis Policy is used in Multiton Inspector

This example illustrates overriding attributes in child policy and how the merged configuration will be used in the child policy for any instance. Any overrides defined in the child policy will be merged with the parent policy. Thus, if attribute1 and attribute2 are overridden in parent policy and attribute2 and attribute3 are overridden in the child policy, the merged configurations are for child policy. This means that attribute1

(defined in parent policy), attribute2 (defined in child policy), and attribute3 (defined in child policy) will be configured on the device.

Parent Policy:

Here we have defined a custom instance by the name **`telnet_parent_instance`** and overridden 2 attributes namely, **`normalize`** and **`encrypted_traffic`** in the custom instance.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": false
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

Child Policy:

This network analysis policy has the aforementioned policy as its base policy. We have overridden attribute **`encrypted_traffic`** from parent policy and also overridden new attribute **`ayt_attack_thresh`**.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "encrypted_traffic": true,
          "ayt_attack_thresh": 1
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  }
}
```

With the above policy JSON, when you deploy the network analysis policy the following merged JSON will be configured on the device.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": true,
          "ayt_attack_thresh": 1
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

This example illustrates details for the custom network analysis policy. The same behavior is also exhibited in the default instance. Also, a similar merging would be done for Singleton inspectors.

Removing all the Inspector Overrides for the Network Analysis Policy

Whenever you want to remove all the overrides for a specific network analysis policy, you can upload an empty JSON. While uploading the overrides, choose the option **Replace inspector overrides**.

Network analysis policy settings and cached changes

When you create a new network analysis policy, it has the same settings as its base policy.

When tailoring a network analysis policy, especially when disabling inspectors, keep in mind that some inspectors and intrusion rules require that traffic first be decoded or preprocessed in a certain way. If you disable a required inspector, the system automatically uses it with its current settings, although the inspector remains disabled in the network analysis policy web interface.

Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

The system caches one network analysis policy per user. While editing a network analysis policy, if you select any menu or other path to another page, your changes stay in the system cache even if you leave the page.

Generate Snort 3 recommendations

Snort 3 rule recommendations

Snort 3 rule recommendations are a security tuning feature that

- automatically tune your intrusion policy with rules specific to the host environment,
- enable additional rules or tune the current rule set by disabling rules for vulnerabilities not present in your network, and
- use the host database to determine rules that apply to your environment.

How rule recommendations work

The management center builds a database of hosts on your network with details such as the IP address, hostname, operating system, services, users, and client applications through passive discovery. Based on this information, the system maps vulnerabilities to each discovered host. The recommendations feature uses this host database to determine the rules that apply to your environment.

In Snort 3, there are four security levels, each corresponding to a specific Talos policy:

- Level 1–Connectivity Over Security
- Level 2–Balanced Security and Connectivity
- Level 3–Security Over Connectivity
- Level 4–Maximum Detection

Check the **Accept Recommendations to Disable Rules** check box to disable rules for vulnerabilities not found on the hosts in your network. Check this option only if you have to trim your rule set because of a high number of alerts, or to improve inspection performance.

For more information, see Secure Firewall recommended rules, on page 2850.

Benefits

- By configuring recommendations, you can tailor your intrusion policy to detect specific types of threats more effectively using rules that are specific to the host environment.
- Recommendations contribute to a more efficient and effective incident response process by reducing false positives and false negatives.

Sample business scenario

A sample business scenario is a network security implementation that demonstrates how a large corporate network uses Snort 3 as its primary intrusion detection and prevention system to enhance incident response capabilities through vulnerability-based rule generation and intrusion policy optimization.

Business scenario details

In this scenario, a large corporate network adopts robust network security measures in a rapidly evolving threat landscape. The security team wants to enhance their incident response capabilities by generating recommendations or rule sets based on the vulnerabilities detected in the host network. This helps to optimize their intrusion policies, thereby safeguarding the network more effectively.

Best practices

- You must have high-quality, accurate host data.

  Because of the passive nature of network discovery, your threat defense devices must be positioned as close as possible to your protected hosts. This allows the threat defense devices to watch network traffic to and from these hosts, giving you accurate data about applications, services, and vulnerabilities on your network.

- Devices should have visibility to east–west traffic flows as well as north–south traffic flows to build an accurate host profile.
- You can create a scheduled task to update recommendations automatically.

Prerequisites for generating recommendations

- Ensure that hosts are present in the system to generate recommendations.
- Protected networks configured for recommendations should map to the hosts present in the system.

Generate Snort 3 recommendations

Generate and apply Snort 3 rule recommendations to optimize your intrusion policy based on your network configuration and security requirements.

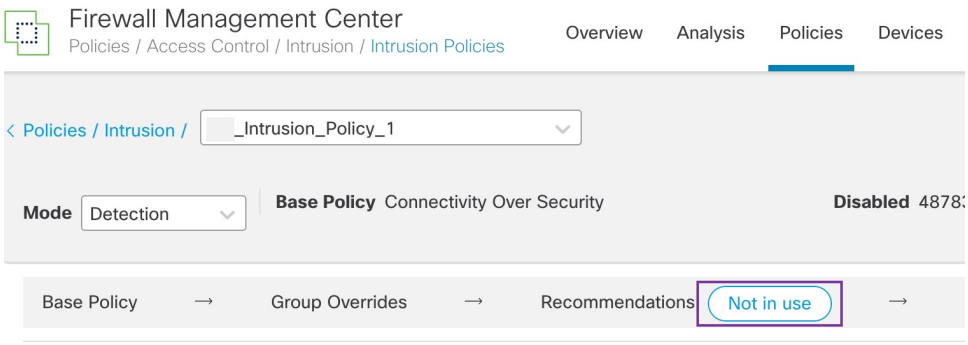
Snort 3 recommendations help you configure intrusion policies with rules that are tailored to your network environment. The system analyzes your protected networks and suggests appropriate security levels and rule configurations.

Procedure

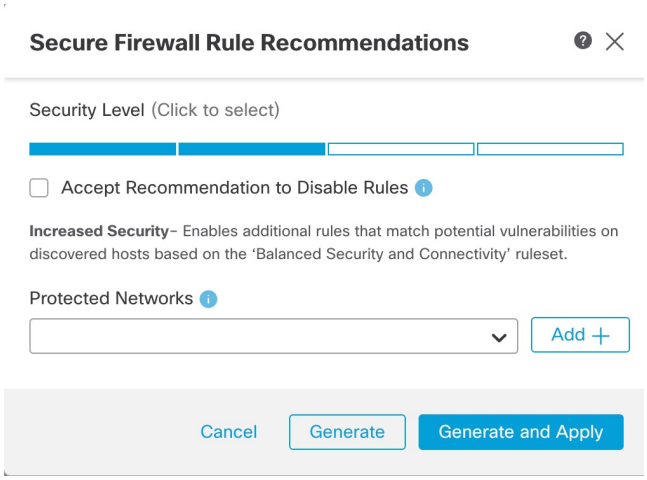
Step 1 Choose **Policies** > **Access Control heading** > **Intrusion**.

Step 2 Click the **Snort 3 Version** button of the corresponding intrusion policy.

Step 3 Click the **Recommendations (Not in Use)** layer to configure the rule recommendations.



In the **Cisco Recommended Rules** window, you can set the security level.



Step 4 Click to select the security level.

Step 5 (Optional) Check the **Accept Recommendation to Disable Rules** check box to disable the rules written for vulnerabilities not found on the hosts in your network.

Use this option, only if you have to trim your rule set because of a high number of alerts or to improve inspection performance.

Step 6 From the **Protected Networks** drop-down list, choose the network objects that must be examined by the recommendations. By default, any IPv4 or IPv6 networks are selected if you do not make a selection.

Click **Add +** to create a new network object of type **Host** or **Network** and click **Save**.

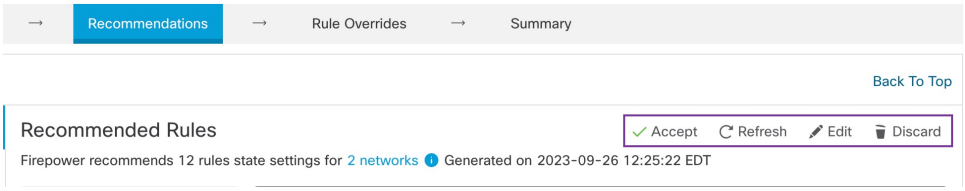
Step 7 Generate and apply recommendations:

- **Generate**—Generates the recommendations for an intrusion policy. This action lists the rules under **Recommended Rules (Not in use)**.
- **Generate and Apply**—Generates and applies the recommendations for an intrusion policy. This action lists the rules under **Recommended Rules (Not in use)**.

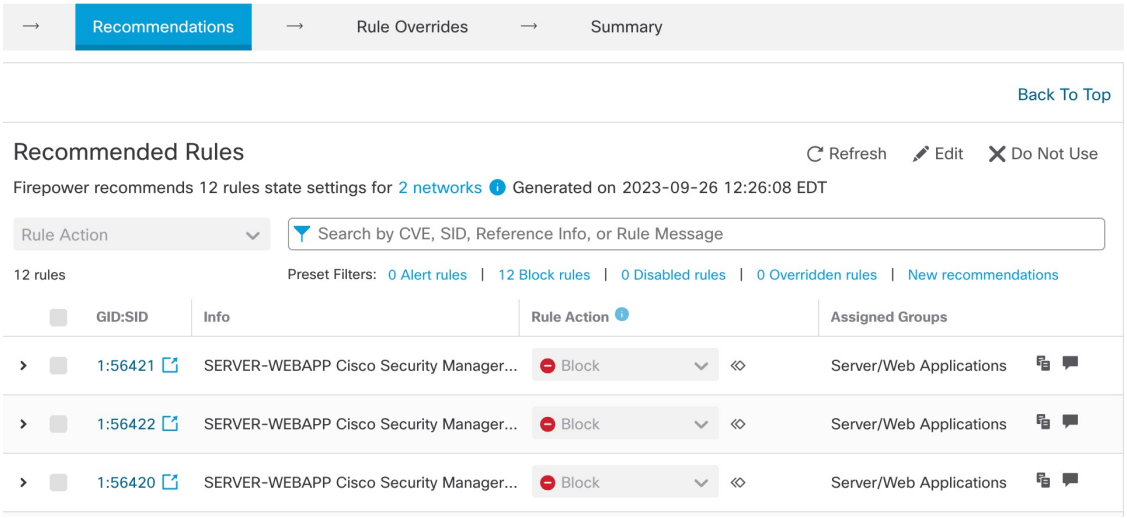
Recommendations are generated successfully. A new recommendation tab appears with all the recommended rules and their corresponding recommended actions. Rule action preset filters are also available for this tab, in addition to new recommendations.

Step 8 Verify the recommendations and then apply them accordingly:

- **Accept**—Applies the previously generated recommendations for an intrusion policy.
- **Refresh**—Regenerates and updates the rule recommendations for an intrusion policy.
- **Edit**—Opens the **Recommendations** dialog box where you can provide the recommendation input values and then generate the recommendations.
- **Discard**—Either reverts or removes the applied recommended rules from the policy; also removes the **Recommendations** tab.



Under **All Rules**, the Recommended Rules section displays the recommended rules.



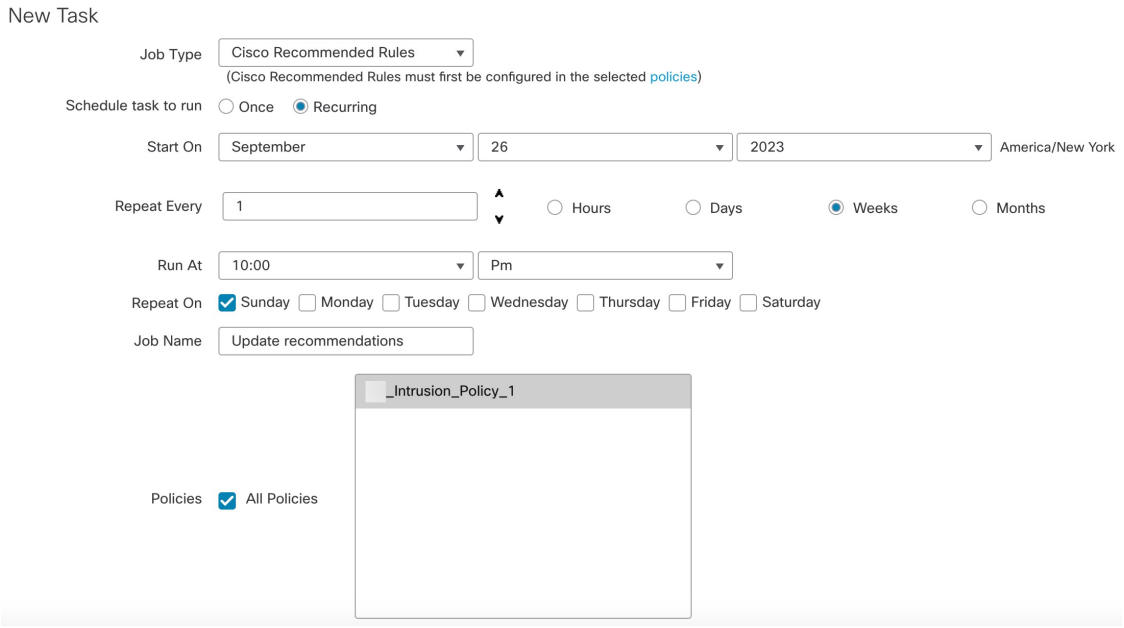
Step 9 To effectively use recommendations, they must be updated periodically. Follow these steps:

a. Choose **System** (⚙) > **Tools** > **Scheduling**.

b. Click **Add Task**.

c. Choose **Cisco Recommended Rules** from the **Job Type** drop-down list.

d. Update the required fields, as needed.



e. Click **Save**.

What to do next

Deploy configuration changes. Refer to Deploy configuration changes, on page 2806.



CHAPTER 102

Advanced Access Control settings for Network Analysis and Intrusion Policies

This chapter describes how to configure advanced settings for network analysis and intrusion policies.

- Advanced access control settings for network analysis and intrusion policies, on page 2917
- Advanced access control settings for network analysis and intrusion policies, on page 2917
- Requirements and prerequisites for advanced access control settings for network analysis and intrusion policies, on page 2918
- Inspection of packets that pass before traffic is identified, on page 2918
- Advanced settings for network analysis policies, on page 2920

Advanced access control settings for network analysis and intrusion policies

Advanced access control settings are intrusion detection and prevention configurations that

- require specific expertise to configure,
- typically require little or no modification, and
- are not common to every deployment.

Version compatibility

Advanced access control settings for network analysis and intrusion policies

Advanced access control settings are intrusion detection and prevention configurations that

- require specific expertise to configure,
- typically require little or no modification, and
- are not common to every deployment.

Version compatibility

Requirements and prerequisites for advanced access control settings for network analysis and intrusion policies

This reference provides the system requirements, software prerequisites, and configuration dependencies needed to implement advanced access control settings for network analysis and intrusion detection policies in Cisco Secure Firewall environments.

Model support

Any.

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Inspection of packets that pass before traffic is identified

Inspection of packets that pass before traffic is identified is a network security process that

- allows a few packets to pass initially for connection establishment and traffic identification,
- enables features like URL filtering, application detection, rate limiting, and Intelligent Application Bypass to function properly, and
- requires explicit configuration in access control policies to inspect, prevent, and generate events for these initial packets.

Configuration requirements

You must explicitly configure your access control policy to inspect these packets, prevent them from reaching their destination, and generate any events.

As soon as the system identifies the access control rule or default action that should handle the connection, the remaining packets in the connection are handled and inspected accordingly.

Best practices for handling packets that pass before traffic identification

The default action specified for an access control policy is NOT applied to these packets.

Use these guidelines to choose a value for the **Intrusion Policy used before Access Control rule is determined** setting in the Advanced settings of the access control policy.

- You can choose a system-created or custom intrusion policy. For example, you can choose **Balanced Security and Connectivity**.
- For performance reasons, unless you have good reason to do otherwise, this setting should match the default action set for your access control policy.
- If your system does NOT perform intrusion inspection (for example, in a discovery-only deployment), select **No Rules Active**. The system will NOT inspect these initial packets, and they will be allowed to pass.
- By default, this setting uses the default variable set. Ensure that this is suitable for your purposes. For information, see Variable sets, on page 3040.
- The network analysis policy associated with the first matching network analysis rule preprocesses traffic for the policy you select. If there are no network analysis rules, or none match, the default network analysis policy is used.

Specify a policy to handle packets that pass before traffic identification

This task explains how to specify the system's behavior for network traffic during the period before traffic identification is complete.

Note This setting is sometimes referred to as the *default intrusion policy*. (This is distinct from the default action for an access control policy.)

Before you begin

Review best practices for these settings. See Best practices for handling packets that pass before traffic identification, on page 2918.

Follow these steps to specify a policy to handle packets that pass before traffic identification:

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to the **Network Analysis** and **Intrusion Policies** section.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Select an intrusion policy from the **Intrusion Policy used before Access Control rule is determined** drop-down list.

If you choose a user-created policy, you can click **Edit** ( ) to edit the policy in a new window. You cannot edit system-provided policies.

Step 3 Optionally, select a different variable set from the **Intrusion Policy Variable Set** drop-down list. You can also select **Edit** (✎) next to the variable set to create and edit variable sets. If you do not change the variable set, the system uses a default set.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Related Topics

Variable sets, on page 3040

Advanced settings for network analysis policies

A *network analysis policy* is a traffic preprocessing configuration that governs how traffic is decoded and preprocessed so that it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.

Advanced configuration concepts

This traffic preprocessing occurs after Security Intelligence matching and traffic decryption, but before intrusion policies inspect packets in detail. By default, the system-provided Balanced Security and Connectivity network analysis policy is the default network analysis policy.

Tip The system-provided Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default. If you are an advanced user with a complex deployment, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones, networks, or VLANs.

To accomplish this, you add custom *network analysis rules* to your access control policy. A network analysis rule is simply a set of configurations and conditions that specifies how you preprocess traffic that matches those qualifications. You create and edit network analysis rules in the advanced options in an existing access control policy. Each rule belongs to only one policy.

Each rule has:

- a set of rule conditions that identifies the specific traffic you want to preprocess
- an associated network analysis policy that you want to use to preprocess traffic that meets all the rules' conditions

When it is time for the system to preprocess traffic, it matches packets to network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rules is preprocessed by the default network analysis policy.

Set the default network analysis policy

Setting a default network analysis policy establishes the baseline preprocessing behavior for network traffic analysis within your access control policy.

You can choose a system- or user-created policy.

Note If you disable a preprocessor but the system needs to evaluate preprocessed packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although it remains disabled in the network analysis policy web interface. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful that you allow the network analysis and intrusion policies examining a single packet to complement each other.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to the Network Analysis and Intrusion Policies section.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 From the **Default Network Analysis Policy** drop-down list, select a default network analysis policy.

If you choose a user-created policy, you can click **Edit** ( ) to edit the policy in a new window. You cannot edit system-provided policies.

Step 3 Click **OK**.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Related Topics

Limitations of custom policies

Network analysis rules

A network analysis rule is a preprocessing configuration component that

- tailors preprocessing configurations to network traffic within access control policy advanced settings,

- matches packets in top-down order by ascending rule number, starting at 1, and
- preprocesses traffic according to the first rule where all conditions match.

Network analysis rule conditions and matching

You can add zone, network, and VLAN tag conditions to a rule. If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion. For example, a rule with a network condition but no zone condition evaluates traffic based on its source or destination IP address, regardless of its ingress or egress interface. Traffic that does not match any network analysis rules is preprocessed by the default network analysis policy.

Network analysis policy rule conditions

Network analysis policy rule conditions are configuration parameters that enable you to fine-tune your network analysis policy to target the users and networks you want to control.

Related Topics

Security zone rule conditions, on page 2202
Network rule conditions, on page 1701
VLAN tag rule conditions, on page 2070

Security zone rule conditions

A security zone rule condition is a network traffic control mechanism that

- segments your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices,
- controls or decrypts traffic by its source and destination security zones, and
- requires all interfaces in a zone to be of the same type (all inline, passive, switched, or routed).

Zone condition traffic matching

If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

All zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Network rule conditions

A network condition is a traffic matching criterion that

- controls or decrypts traffic based on source and destination IP addresses using inner packet headers,
- enables administrators to specify exact IP addresses or blocks for granular policy enforcement, and
- can be built using predefined network objects or by manually entering IP addresses or address blocks.

Network rule condition restrictions

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

VLAN tag rule conditions

VLAN tags rule conditions are filtering mechanisms that

- control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic
- use the innermost VLAN tag to filter VLAN traffic, with the exception of prefilter policy which uses the outermost VLAN tag, and
- apply only to inline sets and do not match traffic on firewall interfaces.

VLAN tag rule conditions configuration details

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note this Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

Configure network analysis rules

Network analysis rules allow you to specify which network analysis policy should preprocess traffic that matches specific conditions, providing granular control over traffic inspection before intrusion detection.

Network analysis rules are configured within the advanced settings of an access control policy. You can create custom rules to apply different network analysis policies based on traffic conditions such as source, destination, ports, or applications.

Before you begin

Follow these steps to configure network analysis rules:

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to the Network Analysis and Intrusion Policies section.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Tip
Click **Network Analysis Policy List** to view and edit existing custom network analysis policies.

Step 2 Next to **Network Analysis Rules**, click the statement that indicates how many custom rules you have.

Step 3 Click **Add Rule**.

Step 4 Configure the rule's conditions by clicking the conditions you want to add.

Step 5 Click **Network Analysis** and choose the **Network Analysis Policy** you want to use to preprocess the traffic matching this rule.

Click **Edit** ( ) to edit a custom policy in a new window. You cannot edit system-provided policies.

Step 6 Click **Add**.

What to do next

- Deploy configuration changes.

Manage network analysis rules

Network analysis rules allow you to control how traffic preprocessing occurs by defining specific conditions and configurations that determine traffic handling methods.

A network analysis rule is simply a set of configurations and conditions that specifies how you preprocess traffic that matches those qualifications. You create and edit network analysis rules in the advanced options in an existing access control policy. Each rule belongs to only one policy.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to the Intrusion and Network Analysis Policies section.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Next to **Network Analysis Rules**, click the statement that indicates how many custom rules you have.

Step 3 Edit your custom rules. You have the following options:

- To edit a rule's conditions, or change the network analysis policy invoked by the rule, click **Edit** ( ) next to the rule.
- To change a rule's order of evaluation, click and drag the rule to the correct location. To select multiple rules, use the Shift and Ctrl keys.
- To delete a rule, click **Delete** ( ) next to the rule.

Tip
Right-click a rule to display a context menu that lets you cut, copy, paste, edit, delete, or add new network analysis rules.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.



PART XXII

Objects and Certificates

- Object Management, on page 2929
- Certificates, on page 3079



CHAPTER 103

Object Management

This chapter describes how to manage reusable objects.

- Objects, on page 2930
- Use the object manager, on page 2932
- AAA servers, on page 2943
- Access lists, on page 2948
- Configure address pools, on page 2954
- Create an application filter, on page 2955
- Configure an AS PATH object, on page 2956
- Configure a BFD template, on page 2956
- Cipher suite lists, on page 2958
- Configure a community list, on page 2959
- DHCP IPv6 pools, on page 2962
- Distinguished names, on page 2962
- DNS server groups, on page 2965
- External attributes, on page 2966
- File lists, on page 2973
- FlexConfig objects, on page 2978
- Geolocation objects, on page 2979
- Interfaces, on page 2980
- Key chains, on page 2980
- Networks, on page 2982
- PKI objects, on page 2986
- Configure a policy list, on page 3012
- Ports, on page 3014
- Prefix lists, on page 3015
- Configure a route map, on page 3017
- Security intelligence, on page 3021
- Sinkhole objects, on page 3033
- Configure an SLA monitor, on page 3034
- Time ranges, on page 3036
- Time zones, on page 3038
- Tunnel zones, on page 3038
- URLs, on page 3038

- Variable sets, on page 3040
- VLAN tags, on page 3057
- VPN, on page 3057

Objects

An object is a configuration element that

- associates a name with a value for increased flexibility and ease of use in the web interface,
- enables reuse of configurations across policies, rules, event searches, reports, and dashboards, and
- can be managed, grouped, and overridden as needed within the system.

Object management and usage

The system uses named objects to simplify configuration and management. Use the object manager to create and manage objects. Many configurations that use objects also allow you to create objects as needed. The object manager enables you to:

- View the policies, settings, and other objects where a network, port, VLAN, or URL object is used; see View objects and their usage, on page 2935.
- Group objects to reference multiple objects with a single configuration; see Object groups, on page 2937.
- Override object values for selected devices; see Object overrides, on page 2939.

After editing an object used in an active policy, you must redeploy the changed configuration for your changes to take effect. You cannot delete an object that is in use by an active policy.

Note The system configures an object on a managed device only when a policy assigned to that device uses the object. If you remove an object from all policies assigned to a given device, the object is also removed from the device configuration on the next deployment, and subsequent changes to the object are not reflected in the device configuration.

Object types

This table lists the objects you can create in the system, and indicates whether each object type can be grouped or configured to allow overrides.

Object type	Groupable?	Allows overrides?
Network	yes	yes
Port	yes	yes
Interface: • Security Zone • Interface Group	no	no

Object type	Groupable?	Allows overrides?
Tunnel Zone	no	no
Application Filter	no	no
VLAN Tag	yes	yes
External Attribute: Security Group Tag (SGT) and Dynamic Object	no	no
URL	yes	yes
Geolocation	no	no
Time Range	no	no
Variable Set	no	no
Security Intelligence: Network, DNS, and URL lists and feeds	no	no
Sinkhole	no	no
File List	no	no
Cipher Suite List	no	no
Distinguished Name	yes	no
Public Key Infrastructure (PKI): • Internal and Trusted CA • Internal and External Certs	yes	no
Key Chain	no	yes
DNS Server Group	no	no
SLA Monitor	no	no
Prefix List: IPv4 and IPv6	no	yes
Route Map	no	yes
Access List: Standard and Extended	no	yes
AS Path	no	yes
Community List	no	yes
Policy List	no	yes
FlexConfig: Text and FlexConfig objects	no	yes

Use the object manager

The object manager provides capabilities for creating and managing objects and object groups, viewing them in paginated lists, filtering results, and refreshing the displayed information.

- The object manager enables creation and management of objects and object groups.
- The object manager displays 20 objects or groups per page. If you have more than 20 of any type of object or group, use the navigation links at the bottom of the page to view additional pages. You can also go to a specific page or click **Refresh** ( ) to refresh your view.
- Access the page using **Objects** > **Object Management**. You can alternatively access the page using **Objects** > **Other FTD Objects**.
- By default, objects and groups are listed alphabetically by name. You can filter objects by name or value.

Import objects

Import objects into the system by uploading a properly formatted comma-separated values file. This task enables bulk addition of supported object types, streamlining object management.

Objects can be quickly imported into the system using a properly formatted CSV file. Up to 1000 objects can be imported in a single operation. The format and required columns vary depending on object type.

Object TYPE	Rules

Individual object	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • DN
	• Both NAME and DN column entries are mandatory to import an entry.
	• You can import individual objects directly into an existing distinguished NAME object group.

Object TYPE	Rules

Network object	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • DESCRIPTION
	  • TYPE
	  • VALUE
	  • LOOKUP
	• To import a host, range, or network object type, provide entries for both NAME and VALUE columns.
	• For an FQDN object, specify "FQDN" in the TYPE column entry, and enter "IPv4," "IPv6," or "IPv4/IPv6" in the LOOKUP column entry.
	• If you leave the LOOKUP column entry for the FQDN object empty, the system saves the object with ipv4_ipv6 as the VALUE.

PORT	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • PROTOCOL
	  • PORT
	  • ICMPCODE
	  • ICMPTYPE
	• The NAME column entry is mandatory.
	• The PORT column entry is mandatory for 'tcp' and 'udp' PROTOCOL types.
	• The ICMPCODE and ICMPTYPE column entries are mandatory for 'icmp' and 'icmp6' PROTOCOL types.

Object TYPE	Rules

URL	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • DESCRIPTION
	  • URL
	• You must provide entries for the NAME and URL columns to import an object.

VLAN TAG	• The column header must be mentioned in capital letters.
	• The file must have these columns headers:
	  • NAME
	  • DESCRIPTION
	  • TAG
	• The NAME and TAG column entries are mandatory to import an entry.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose one of these object types from the left pane:

- **Distinguished Name** > **Individual Objects**
- **Network Object**
- **PORT**
- **URL**
- **VLAN TAG**

Step 3 Choose **Import Object** from the **Add [Object TYPE]** drop-down list.

Note
If you have selected **Individual Objects** in the previous step, click **Import**.

Step 4 Click **Browse**.

Step 5 Locate and select the comma-separated file on your system.

Step 6 Click **Open**.

Note

While importing **Distinguished Name** objects, you can optionally check the **Add imported Distinguished Name objects to the below object group** check box and select the group name from the drop-down box to import the objects directly to an existing distinguished name object group.

Step 7 Click **Import**.

Edit objects

Edit network objects and groups to update their settings or manage variable sets as required for your deployment.

See caveats related to network objects and groups at Networks, on page 2982.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose an object type from the list; see Objects, on page 2930.

Step 3 Click **Edit** ( ) next to the object you want to edit.

If **View** ( ) appears instead, the object belongs to an ancestor domain and has been configured not to allow overrides, or you do not have permission to modify the object.

Step 4 Modify the object settings as desired.

Step 5 If you are editing a variable set, manage the variables in the set; see Manage variables, on page 3054.

Step 6 For objects that can be configured to allow overrides:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941. You can change this setting only for objects that belong to the current domain.
- If you want to add override values to this object, expand the Override section and click **Add**; see Add object overrides, on page 2942.

Step 7 Click **Save**.

Step 8 If you are editing a variable set, and that set is in use by an access control policy, click **Yes** to confirm that you want to save your changes.

What to do next

- If an active policy references your object, deploy configuration changes.

View objects and their usage

This task enables you to view where objects are used within your system, helping you understand object dependencies and usage across policies and settings.

You can view usage details of objects on the Object Management page. Cloud-Delivered Firewall Management Center provides this functionality for many object types. However, some object types are not supported.

Before you begin

No prerequisites are required to view object usage details.

Follow these steps to view objects and their usage:

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose one of the following supported object types:

- Access List > Extended
- Access List > Standard
- AS Path
- Community List
- Interface
- Network
- Policy List
- Port
- Prefix List > IPv4 Prefix List
- Prefix List > IPv6 Prefix List
- Route Map
- SLA Monitor
- URL
- VLAN Tag

Step 3 Click the **Find Usage** () icon next to the object.

The Object Usage window displays a list of all the policies, objects, and other settings where the object is in use. Click any of the listed items to know more about the object usage. For policies and some other settings where the object is used, you can click the corresponding links to visit the respective UI pages.

After completing these steps, you will be able to see where each object is used within your system, allowing you to manage and audit object usage efficiently.

Filter objects or object groups

Use this task when working with large configurations or when you need to identify and manage unused objects or object groups.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Enter your filter criteria in the **Filter** field.

The page updates as you type to display matching items.

You can use these wildcards:

- The asterisk (*) matches zero or more occurrences of a character.
- The caret (^) matches content at the beginning of a string.
- The dollar sign ($) matches content at the end of a string.

Step 3 To identify and review unused objects or object groups, check the **Show Unused Object** check box.

Note

- If an object is a part of an unused object group, the object is considered as used. However, the unused object group is displayed when the **Show Unused Object** check box is checked.
- The **Show Unused Object** check box is available only for network, port, URL, and VLAN tag object types.

Object groups

An object group is a configuration grouping that

- enables referencing multiple objects with a single configuration,
- allows objects and object groups to be used interchangeably in the web interface, and
- supports grouping of network, port, VLAN tag, URL, and PKI objects.

Object group usage and management

The system allows you to use objects and object groups interchangeably in the web interface. For example, anywhere you would use a port object, you can also use a port object group.

Network object groups can be nested. You can add a network object group to another network object group, up to 10 levels deep.

Objects and object groups of the same type cannot have the same name.

Edit an object group used in a policy, such as a network object group in an access control policy. Then, re-deploy the configuration for your changes to take effect.

Deleting a group does not delete the objects in the group; only their association with each other. Additionally, you cannot delete a group that is in use in an active policy. For example, you cannot delete a VLAN tag group that is used in a VLAN condition in a saved access control policy.

Object references and literal entries in object groups

Object groups can contain either object references or literal entries.

- Object references are links to existing Host/Network Objects. When the original object is modified (for example, an IP address change), all object groups containing references to that object are automatically updated to reflect the change.
- Literal entries are manually created entries within an object group. Unlike object references, literal entries are independent configuration entities and are not dynamically linked to existing Host/Network Objects. Therefore, changes made to a Host/Network Object do not propagate to object groups containing literal entries.

For example, an administrator may manually add a literal entry with the same name and IP address as an existing Host/Network Object. Although the entries appear identical, the literal entry remains independent of the original object. If the original Host/Network Object is modified, only object groups containing a reference to that object are updated.

Modifying a Host/Network Object updates only those object groups that contain a reference to that object. Object groups containing literal entries continue to display their original configuration unless the literal entries are manually modified.

Group reusable objects

Group reusable objects to simplify configuration and management in Secure Firewall Management Center.

- Organize object types such as Network, Port, URL, VLAN Tag, Distinguished Name, and PKI into groups for easier policy application.

This task is relevant when you need to manage multiple reusable objects and apply them efficiently in firewall policies.

Grouping objects allows for streamlined configuration and reduces manual effort when deploying or updating policies.

- Use this task when creating or editing object groups for policy deployment in Secure Firewall Management Center.

Before you begin

No explicit prerequisites are stated. Ensure you have access to Secure Firewall Management Center and the necessary permissions to manage objects.

- Identify the object types you want to group before starting the task.

Follow these steps to group reusable objects:

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 If the object type you want to group is **Network**, **Port**, **URL**, or **VLAN Tag**:

a) Choose the object type from the list of object types.

b) Choose **Add Group** from the **Add [Object Type]** drop-down list.

Step 3 If the object type you want to group is **Distinguished Name**:

a) Expand the **Distinguished Name** node.
b) Choose **Object Groups**.
c) Click **Add Distinguished Name Group**.

Step 4 If the object type you want to group is **PKI**:

a) Expand the **PKI** node.
b) Choose one of the following:

- **Internal CA Groups**
- **Trusted CA Groups**
- **Internal Cert Groups**
- **External Cert Groups**

c) Click **Add [Object Type] Group**.

Step 5 Enter a unique **Name**.

Step 6 Choose one or more objects from the list, and click **Add**.

You can also:

- Use the filter field **Search** ( ) to search for existing objects to include, which updates as you type to display matching items. Click **Reload** ( ) above the search field or click **Clear** ( ) in the search field to clear the search string.
- Click **Add** ( ) to create objects on the fly if no existing objects meet your needs.

Step 7 Optionally for **Network**, **Port**, **URL**, and **VLAN Tag** groups:

- Enter a **Description**.
- Check the **Allow Overrides** check box to allow overrides for this object group; see Allow object overrides, on page 2941.

Step 8 Click **Save**.

After completing these steps, the object group is created and available for use in firewall policies. You can reference the group in active policies and deploy configuration changes as needed.

What to do next

- If an active policy references your object group, deploy configuration changes; see Deploy configuration changes, on page 202.

Object overrides

An object override is a configuration mechanism that

- enables the definition of alternate values for objects, which the system applies to specified devices,

- lets administrators create a smaller set of shared policies that can be used across devices without losing the flexibility to modify policies for individual devices, and
- allows modifications to an object for devices or domains that require different definitions.

Object override usage and supported object types

Administrators can create an object whose definition works for most devices and use overrides to specify modifications for those devices needing different definitions. You may also need to override an object for all devices. This approach allows you to create a single policy that applies everywhere. Object overrides allow you to use fewer shared policies and customize them for individual devices when necessary.

You can target overrides to a specific domain. The override value applies to all devices in that domain unless you set a device-specific override. The object manager enables selection of objects that can be overridden and allows listing device-level or domain-level overrides for each object.

Supported object types for overrides include:

- Network
- Port
- VLAN tag
- URL
- SLA Monitor
- Prefix List
- Route Map
- Access List
- AS Path
- Community List
- Policy List
- Cert Enrollment (PKI)
- Key Chain

In the object manager, the **Override** column flags objects that support overrides:

- Green checkmark: Overrides can be created, and none have been added yet.
- Red X: Overrides are not supported for the object.

- Number: Indicates the count of overrides added for the object (for example, "2" means two overrides have been defined).

Object override example

For example, to deny ICMP traffic to different departments, each connected to a different network, you can define an access control policy that includes a network object called Departmental Network. By enabling overrides for this object, it's possible to create policy variations on each relevant device that specify the actual network associated with that device.

Manage object overrides

Object overrides are used when you need to adjust the behavior of reusable objects for specific use cases, such as applying different values in different domains or policies. Use this task when you want to add, allow, delete, or edit object overrides in the Secure Firewall Management Center.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose from the list of object types; see Objects, on page 2930.

Step 3 Click **Edit** ( ) next to the object you want to edit.

If **View** ( ) appears instead, the object belongs to an ancestor domain and has been configured not to allow overrides, or you do not have permission to modify the object.

Step 4 Manage the object overrides:

- Add—Add object overrides; see Add object overrides, on page 2942.
- Allow—Allow object overrides; see Allow object overrides, on page 2941.
- Edit—Edit object overrides; see Edit an object override, on page 2942.
- Delete—In the object editor, click **Delete** ( ) next to the override you want to remove.

Step 5 Click **Save**.

Allow object overrides

Enable the Allow Overrides option for an object so that you can specify override values as required for different use cases.

Procedure

Step 1 In the object editor, check the **Allow Overrides** check box.

Step 2 Click **Save** to apply your changes.

What to do next

Add object override values; see Add object overrides, on page 2942.

Add object overrides

Add or adjust object overrides so that network objects or groups can have different values assigned in specific domains or to specific devices, making configuration more flexible without modifying the base object.

Use this task when you need to assign different settings or values to the same object across multiple domains or devices—for example, when different parts of your organization require different configurations for the same network object.

For caveats when working with network objects or groups, see Networks, on page 2982.

Before you begin

Ensure that object overrides are allowed in your environment. See Allow object overrides, on page 2941.

Procedure

Step 1 In the object editor, expand the **Override** section.

Step 2 Click **Add**.

Step 3 On **Targets**, choose domains or devices in the **Available Devices and Domains** list and click **Add**.

Step 4 On the Override tab, enter a **Name**.

Step 5 Optionally, enter a **Description**.

Step 6 Enter an override value.

Example:

```
For a network object, enter a network value.
```

Step 7 Click **Add**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Edit an object override

This task enables you to edit the description and value of an existing object override to update its configuration as needed.

You can modify the description and the value of an existing override, but you cannot modify the existing target list. Add a new override with new targets to replace the existing override.

For caveats when working with network objects or groups, see Networks, on page 2982..

Procedure

Step 1 In the object editor, expand the **Override** section.

Step 2 Click **Edit** ( ) next to the override you want to modify.

Step 3 Optionally, modify the **Description**.

Step 4 Modify the override value.

Step 5 Click **Save** to save the override.

Step 6 Click **Save** to save the object.

What to do next

- If an active policy references your object, deploy configuration changes.

AAA servers

An AAA server is a network authentication, authorization, and accounting component that

- enables centralized management of user credentials
- supports policy enforcement, and
- provides accounting for user activities.

Add a RADIUS server group

RADIUS Server Group objects contain one or more references to RADIUS servers. These servers are used to authenticate users logging in through Remote Access VPN connections.

You can apply these objects to compatible Firewall Threat Defense devices.

Before you begin

Note You cannot override RADIUS Server Group Objects.

Procedure

Step 1 Select **Objects** > **Object Management** > **AAA Server** > **RADIUS Server Group**.

All currently configured RADIUS Server Group objects will be listed. Use the filter to narrow down the list.

Step 2 Choose and edit a listed RADIUS Server Group object, or add a new one.

See Configure RADIUS server options, on page 2945 and RADIUS server group configuration options, on page 2944 to configure this object.

Step 3 Click **Save**.

RADIUS server group configuration options

The configuration options for a RADIUS server group enable integration with identity policies and VPN authentication. This topic lists each option and describes its purpose, required settings, and usage notes.

Navigation path

Objects > **Object Management** > **AAA Server** > **RADIUS Server Group**. Choose and edit a configured RADIUS Server Group object or add a new one.

Fields

- **Name** and **Description**—Enter a name and optionally, a description to identify this RADIUS Server Group object.
- **Group Accounting Mode**—Defines how accounting messages are sent to the RADIUS servers. Choose one of these:
  - **Single**—Sends accounting messages to a single server (default).
  - **Multiple**—Sends accounting messages to all servers in the group simultaneously.
- **Retry Interval**—Sets the interval (1–10 seconds) between attempts to contact RADIUS servers.
- **Realms**(Optional)—Specify the Active Directory (AD) realm associated with this RADIUS server group. The selected realm is referenced in identity policies to determine VPN authentication identity sources for relevant traffic flows. This realm effectively provides a bridge from the identity policy to this RADIUS server group. If no realm is associated with this RADIUS server group, the RADIUS server group cannot be reached to determine the VPN authentication identity source for a traffic flow in an identity policy.

Note This field is mandatory if you use remote access VPN with User Identity and RADIUS as the identity source.

- **Enable authorize only**—If the group is used only for authorization or accounting (not authentication), enable authorize-only mode. This mode omits the RADIUS server password from Access-Request messages; the configured password for individual RADIUS servers is ignored.
- **Enable interim account update** and **Interval**—Enables the generation of interim accounting update messages to notify the RADIUS server of newly assigned IP addresses. The Interval option sets the period (1 to 120 hours; default is 24 hours) between updates.
- **Enable Dynamic Authorization** and **Port**— Enables the RADIUS dynamic authorization or change of authorization (CoA) services for this RADIUS server group. Specify the listening port for RADIUS CoA requests in the **Port** field. The valid range is 1024 to 65535 and the default value is 1700. Once defined, the corresponding RADIUS server group will be registered for CoA notification and it listens to the port for the CoA policy updates from the Cisco Identity Services Engine (ISE).

- **Merge Downloadable ACL with Cisco AV Pair ACL**—Enables merging a downloadable access control list (dACL) with a Cisco attribute-value (AV) pair ACL.

  A downloadable ACL defines and updates access control lists in Cisco ISE. Administrators can then download the updated ACLs to all applicable controllers. For more information about using dACLs in Cisco ISE, see the chapter on Segmentation, section on authorization policies, in the *Cisco ISE Administrator Guide*.

  A Cisco AV pair ACL can be utilized to define specific authentication, authorization, and accounting elements for each individual session. For more information about using dACLs in Cisco ISE, see the chapter on Segmentation, section on authorization profile settings, in the *Cisco ISE Administrator Guide*.

  If you select **Merge Downloadable ACL with Cisco AV Pair ACL**, you can choose the following options:

  - **After Cisco AV Pair ACL**—Place downloadable ACL entries after AV pair entries.
  - **Before Cisco AV Pair ACL**—Place downloadable ACL entries before AV pair entries.

- **RADIUS Servers**—For server-specific configuration options, see Configure RADIUS server options, on page 2945.

Related Topics

Add a RADIUS server group, on page 2943

Configure RADIUS server options

Use this topic to configure a RADIUS server, including authentication, accounting, connectivity, and security settings.

Navigation path

Objects > **Object Management** > **AAA Server** > **RADIUS Server Group**. Choose and edit a listed RADIUS Server Group object, or add a new one. In the RADIUS Server Group dialog, choose and edit a RADIUS Server or add a new server.

Fields

- **IP Address/Hostname**—Specify the hostname or IP address (IPv4 or IPv6) of the RADIUS server that receives authentication requests. Only one server may be selected. To add servers, add more RADIUS Server entries to the RADIUS Server Group list.
- **RADIUS Server-Enabled Message Authenticator**—Message authenticators safeguard server-to-client communication and are required for a secure connection between your RADIUS server and firewall devices. Disabling message authenticators may expose your firewalls to potential attacks.

  Your RADIUS server must support and be configured for Message-Authenticator. Requires Firewall Threat Defense Version 7.2.10+, 7.4.3+, or 7.6.1+. For the Firepower 4100/9300, this feature may require an FXOS upgrade; for minimum builds, see *Cisco Secure Firewall Threat Defense Compatibility Guide*.
- **Authentication Port**—The port on which RADIUS authentication and authorization are performed. The default is 1812.
- **Key** and **Confirm Key**—Enter the shared secret to encrypt data between the managed device (client) and the RADIUS server.

The key is a case-sensitive, alphanumeric string of up to 127 characters. Special characters are permitted.

Enter the same key in this field and on the RADIUS server. Enter the shared secret again in the Confirm field.

- **Accounting Port**—Enter the port number for RADIUS accounting. The default port is 1813.
- **Timeout**—Specify the session timeout for authentication.

  Set the timeout value to 60 seconds or more for RADIUS two-factor authentication. The default timeout value is 10 seconds.
- **Connect Using** —Select the method to connect the device to a RADIUS server using either a route lookup or a specific interface.
  - Click the **Routing** radio button to use the routing table.
  - Click the **Specific Interface** radio button and choose a security zone/interface group or the Management interface (the default) from the drop-down list. If you want to use Management, you must choose it specifically; it is not available when using a route lookup. You cannot specify any other management-only interface as the RADIUS source.You can also choose a loopback interface group.
- **Redirect ACL**—Select the redirect ACL from the list or add a new one.

  Use the ACL name from the device to control redirected traffic. The Redirect ACL name here must be the same as the *redirect-ACL* name in ISE server. When you configure the ACL object, ensure that you select Block action for ISE and DNS servers, and Allow action for the rest of the servers.

Related Topics

Add a RADIUS server group, on page 2943
RADIUS server group configuration options, on page 2944

Add a single sign-on server

Add a single sign-on (SSO) server to integrate SAML-based authentication for remote access VPN users. This task enables secure user authentication through your identity provider.

Use this task when you need to configure SAML single sign-on for remote access VPN, allowing users to authenticate using an external identity provider. You typically need to perform this configuration when you integrate with SAML IdPs, such as Microsoft Azure.

Before you begin

Obtain the following from your SAML identity provider:

- Identity Provider Entity ID URL
- Sign-in URL
- Sign-out URL
- Identity provider certificate and enroll the certificate in Firewall Threat Defense using the Cloud-Delivered Firewall Management Center web interface (**Devices** > **Certificates**)

For more information, see Configuring a SAML Single Sign-On Authentication, on page 1908.

Follow these steps to add a single sign-on server:

Procedure

Step 1 Choose **Objects** > **Object Management** > **AAA Server** > **Single Sign-on Server**.

Step 2 Click **Add Single Sign-on Server** and provide the following details:

Field	Description
Name	Enter the name for the SAML single sign-on server object.
Identity Provider Entity ID	Enter the URL that the SAML IdP defines to uniquely identify a service provider. Enter the URL for a metadata XML page that describes how the SAML Issuer responds to requests.
SSO URL	Enter the URL for signing in to the SAML identity provider server.
Logout URL	Enter the URL for signing out of the SAML identity provider server.
Base URL	Enter the URL that redirects users back to Firewall Threat Defense after the identity provider authenticates them. This should be the access interface URL configured for the Firewall Threat Defense remote access VPN.
Identity Provider Certificate	Enter the IdP certificate enrolled into the Firewall Threat Defense to verify the signed messages. Select an identity provider certificate from the list or click **Add** to create a new certificate enrollment object. For more information, see Manage Firewall Threat Defense certificates, on page 3080. You must enroll all of the Microsoft Azure registered application CA certificates as Trustpoints on the Firewall Threat Defense. The Microsoft Azure SAML identity provider is configured on Firewall Threat Defense for the initial application. All connection profiles are mapped to the configured MS Azure SAML identity provider. For each of the MS Azure applications (other than the default), you can choose the required trustpoint(CA certificate) in the connection profile configuration of the remote access VPN. For details, see Configure AAA Settings for Remote Access VPN, on page 1835.
Service Provider Certificate	Firewall Threat Defense certificate, which will be used to sign the requests and build circle of trust with IdP. If you have not enrolled internal Firewall Threat Defense certificates, click + to add and enroll a certificate. For more information, see Manage Firewall Threat Defense certificates, on page 3080.
Request Signature	Select the encryption algorithm to sign the SAML single sign-on requests. The signatures are listed from weakest to strongest: SHA1, SHA256, SHA384, SHA512. Select None to disable encryption.

Field	Description
Request Timeout	Specify how long users can complete the single sign-on request (SAML assertion validity duration). The SAML IdP has two time outs: *NotBefore* and *NotOnOrAfter*. The Firewall Threat Defense validates if its current time is within the time range of (lower limit) *NotBefore* and (upper limit) the smaller of *NotBefore* plus *timeout* and *NotOnOrAfter*. Thus, if you set a timeout longer than the IdP's *NotOnOrAfter* timeout, the specified timeout is ignored and the *NotOnOrAfter* timeout is selected. If the sum of the specified timeout and the *NotBefore* timeout is less than the *NotOnOrAfter* time, Firewall Threat Defense timeout overrides the timeout. The timeout range is 1-7200 seconds; the default is 300 seconds.
Enable IdP only accessible on Internal Network	Select this option if the SAML IdP resides on the internal network. Firewall Threat Defense acts as a gateway and establishes communication between the users and IdP using an anonymous webvpn session.
Request IdP re-authentication on Login	Select this option to authenticate user at each login even if the previous IdP session is valid.
Allow Overrides	Select this check box to allow overrides for this single sign-on server object.

Step 3 Click **Save**.

The single sign-on server is added and configured for SAML authentication. Users can now authenticate to the remote access VPN using the specified identity provider.

Access lists

An access list is a traffic selection object that

- identifies and selects traffic to which a service will apply
- is used when configuring features such as route maps for Firewall Threat Defense devices, and
- distinguishes between allowed and blocked traffic, where blocked traffic is excluded from the service but not necessarily dropped altogether.

Types of access lists

You can configure these types of ACL:

- Extended: Identifies traffic based on source and destination address and ports. Supports IPv4 and IPv6 addresses, which you can mix in a given rule.
- Standard: Identifies traffic based on destination address only. Supports IPv4 only.

Each access list is composed of one or more access control entries (ACEs), or rules. The order of ACEs is important.

1. When the ACL is evaluated, a packet is tested against each ACE in the order listed.

2. After a match is found, no more ACEs are checked.
3. Place more specific rules at the top of an ACL for correct behavior.

Place the most specific rules at the top of the ACL so it works as expected

Configure an extended ACL object

Configure extended ACL objects to define traffic classification rules based on source and destination addresses, protocols, ports, application groups, or IPv6 traffic. This enables granular control over network traffic for security and policy enforcement.

Use extended ACL objects when you want to match traffic based on source and destination addresses, protocol and port, application group, or if the traffic is IPv6. These objects are useful for policy-based routing and direct internet access scenarios, and provide flexibility in defining access control entries for various network requirements.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Access List** > **Extended** from the table of contents.

Step 2 Do one of the following:

- Click **Add Extended Access List** to create a new object.
- Click **Edit** ( ) to edit an existing object.

Step 3 In the **New Extended Access List Object** dialog box, enter a name for the object (no spaces allowed) and configure the access control entries:

a) Do one of the following:

- Click **Add** to create a new entry.
- Click **Edit** ( ) to edit an existing entry.

b) Select the **Action**, whether to Allow (match) or Block (not match) the traffic criteria.

Note
The **Logging**, **Log Level**, and **Log Interval** options are used for access rules only (ACLs attached to interfaces or applied globally). Because ACL objects are not used for access rules, leave these values at their defaults.

c) Configure the source and destination addresses on the **Network** tab using any of the following techniques:

- Select the desired network objects or groups from the Available list and click **Add to Source** or **Add to Destination**. You can create new objects by clicking the + button above the list. You can mix IPv4 and IPv6 addresses.
- Type an address in the edit box below the source or destination list and click **Add**. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), or a subnet (in 10.100.10.0/24 or 10.100.10.0 255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60).

d) Click the **Port** tab and configure the service using any of the following techniques.

- Select the desired port objects from the Available list and click **Add to Source** or **Add to Destination**. You can create new objects by clicking the + button above the list. The object can specify TCP/UDP ports, ICMP/ICMPv6 message types, or other protocols (including "any"). However, the source port, which you typically would leave empty, accepts TCP/UDP only. You cannot select port groups.

  For TCP/UDP, note that you must use the same protocol in both the source and destination fields, if you specify both. For example, you cannot specify a UDP source port and a TCP destination port.

- Type or select a port or protocol in the edit box below the source or destination list and click **Add**.

Note
To get an entry that applies to all IP traffic, select a destination port object that specifies "all" protocols.

e) Click the **Application** tab and choose the applications that are to be grouped for the direct internet access policy.

Important
- You cannot configure applications for cluster devices. Hence, this tab is not applicable for cluster devices.
- Use extended ACL with applications only in policy-based routing. Do not use it in other policies as its behavior is unknown and not supported. Ensure migration of the realm/ISE configuration for policy-based routing that uses User Identity and SGT in extended ACL.

Note
- The **Available Applications** list displays a fixed set of pre-defined applications. This list is a subset of the applications that are available on the Access Control policy as only they can be detected by their first packet (FQDN end-points resolved to IP addresses and port). The application definitions are updated through the VDB updates and are pushed to Firewall Threat Defense during subsequent deployments.
- User-defined custom applications or group of applications are not supported.
- Currently, Cloud-Delivered Firewall Management Center neither supports user-defined custom applications or group of applications nor allows you to modify the pre-defined applications list.
- You can use the filter options provided under the **Application Filters** to refine this list.

f) Click the **Users** tab and choose the users, user groups, or both that are to be classified for the Policy Based Routing (PBR).

Important
Use extended ACL with users, user groups, or both only for Policy Based Routing. Do not use it in other policies as its behavior is unknown and not supported.

Note
- The **Available Realms** list displays the configured active directory/LDAP realms. For information on creation of realm and managing them, see Create an LDAP realm or an Active Directory realm and realm directory, on page 2325 and Manage a realm, on page 2341 respectively.

  Note
  Local realms and Azure AD realms are not supported.

- The **Available Users** list displays the downloaded users and user groups of the selected AD/LDAP realms. To download the users, user groups, or both, navigate to **Integrations** > **Other Integrations** > **Realms**, and then click **Download** against the relevant active directory/LDAP realms.

  Note
  Threat defense can support a maximum of 512 user groups and 64000 user-IP mappings.

- The user to IP mapping and user group membership information are updated and pushed to the Firewall Threat Defense from the Cloud-Delivered Firewall Management Center during the user login or logouts, and changes in the group memberships.

g) Click the **Security Group Tag** tab and choose the source SGT tags to be classified for the direct internet access policy.

Important
Use extended ACL with SGTs only for Policy Based Routing. Do not use it in other policies as its behavior is unknown and not supported.

Note
- The **Available Security Group Tags** list displays the configured security group tags. You can choose to use the ISE SGTs or create custom SGTs.
- To use ISE SGTs, ensure that ISE is integrated with Cloud-Delivered Firewall Management Center with **Session Directory Topic** and SXP Topic subscribed.

  Note
  The supported ISE versions are 3.2, 3,1, 3.0, and 2.7 patch 2+.

- For information on creating custom SGTs, see Create a security group tag object, on page 2972.

Important
You can configure **Users** and **Security Group Tag** for Secure Firewall 4200 devices only.

h) Select the required application, and click **Add to Rule**.

Note
- Do not configure destination networks and applications in the extended ACL object.
- The selected applications (Network Service objects) in each of the access control entries, form a Network Service Group (NSG) and this group is deployed on the Firewall Threat Defense. The NSG is used in direct internet access to classify traffic based on the match with the selected application group.

i) Click **Add** to add the entry to the object.

Troubleshooting Tips

If necessary, click and drag the entry to move it up or down in the rule order to the desired location. Repeat the process to create or edit additional entries in the object.

Step 4 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 5 Click **Save**.

Configure a service access object

This task enables you to configure a service access object, which defines the conditions that traffic must meet to access a service, such as remote access VPN on the Firewall Threat Defense device.

Service access objects allow you to manage access based on geolocations and ensure that only traffic from approved regions can access specified services.

A service access object defines conditions as multiple rules to be executed in order. Each service access rule has an Allow or Deny action and a set of match criteria such as country, continent, or user-defined geolocation objects.

These rules are used to manage access to services, such as remote access VPN, and enforce security policies based on the source location of the traffic. If the traffic does not match any rules, the default action is enforced.

Use this task when you need to enforce or restrict access based on geographic or custom geolocation objects.

Before you begin

- Configure geolocation objects. For more information, see Geolocation objects, on page 2979.
- Configure a remote access VPN policy. For more information, see Create a New Remote Access VPN Policy, on page 1822.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Access List** > **Service Access**.

Step 2 Click **Add Service Access Object** to create a new object.

Step 3 In the **Add Service Access Object** dialog box, configure the following parameters:

a) In the **Name** field, enter a name for the object.

b) Click **Add Rule** to add a rule for the service access object. In the **Add Service Access Rule** dialog box, configure the following parameters:

1. Choose an action as **Allow** or **Deny** from the drop-down list.
2. From **Available Countries**, select countries, continents, or user-defined geolocation objects and move them to the **Selected Geolocation** list.
3. Click **Add**.

After you create a service access rule, a sequence number is assigned to the rule. This number determines the execution order of the rule. You cannot reorder these rules.

Step 4 From the **Default Action** drop-down list, choose **Allow All Countries** or **Deny All Countries**.

Step 5 (Optional) Check the **Allow Overrides** check box and click + to configure overrides for the service access object for devices.

Note

When you define an override for a device or a domain, whenever the system configures this object on the device or domain, it uses the override value instead of the value defined in the original object.

Step 6 (Optional) In the **Add Service Access Override** dialog box, configure these parameters:

a) From the **Target Devices and Domains** drop-down list, choose the required devices or domains for the override.

b) Verify if you require the existing rules. Otherwise, delete them.

c) Click **Add Rule** to create service access rules for this object override.

1. Choose an action as **Allow** or **Deny** from the drop-down list.
2. From **Available Countries**, select countries, continents, or user-defined geolocation objects and move them to the **Selected Geolocation** list.
3. Click **Add**.

d) From the **Default Action** drop-down list, choose **Allow All Countries** or **Deny All Countries**.

e) Click **Add**.

Step 7 Click **Save**.

What to do next

Configure the Service Access object in the remote access VPN policy. For more information, see Configure Access Interfaces for Remote Access VPN, on page 1851.

Configure a standard ACL object

Use standard ACL objects when you only need to match traffic based on destination IPv4 address. Extended ACLs allow more granularity if needed.

Procedure

Step 1 Select **Objects** > **Object Management** > **Access List** > **Standard** from the table of contents.

Step 2 Do one of the following:

- Click **Add Standard Access List** to create a new object.
- Click **Edit** (✎) to edit an existing object.

Step 3 In the **New Standard Access List Object** dialog box, enter a name for the object (no spaces allowed), and configure the access control entries:

a) Do one of the following:

- Click **Add** to create a new entry.
- Click **Edit** (✎) to edit an existing entry.

b) For each access control entry, configure the following properties:

- **Action**—Whether to allow (match) or block (not match) the traffic criteria.
- **Network**—Add the IPv4 network objects or groups that identify the destination of the traffic.

c) Click **Add** to add the entry to the object.
d) If necessary, click and drag the entry to move it up or down in the rule order to the desired location.

Repeat the process to create or edit additional entries in the object.

Step 4 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 5 Click **Save**.

Configure address pools

You can configure IP address pools for both IPv4 and IPv6 that can be used for clustering or for VPN remote access profiles. For clustering in Individual interface mode, you can also configure MAC address pools.

Procedure

Step 1 Select **Objects** > **Object Management** > **Address Pools**.

Step 2 Click **IPv4 Pools** and then **Add IPv4 Pools**, and configure the following fields.

- **Name**—Enter the name of the address pool. It can be up to 64 characters.
- **Description**—Add an optional description for this pool.
- **IP Address**—Enter a range of addresses available in the pool. Use dotted decimal notation and a dash between the beginning and the end address, for example: 10.10.147.100-10.10.147.177.
- **Mask**—Identifies the subnet on which this IP address pool resides.
- **Allow Overrides**—Check this check box to enable object overrides. Click the expand arrow to show the **Overrides** table. You can add a new override by clicking **Add**. See Object overrides, on page 2939 for more information.

Step 3 Click **Save**.

Step 4 Click **IPv6 Pools** and then **Add IPv6 Pools**, and configure the following fields.

- **Name**—Enter the name of the address pool. It can be up to 64 characters
- **Description**—Add an optional description for this pool.
- **IPv6 Address**—Enter the first IP address available in the configured pool and the prefix length in bits. For example: 2001:DB8::1/64.
- **Number of Addresses**—Identifies the number of IPv6 addresses, starting at the Starting IP Address, that are in the pool.

- **Allow Overrides**—Check this check box to enable overrides. Click the expand arrow to show the **Overrides** table. You can add a new override by clicking **Add**. See Object overrides, on page 2939 for more information.

Step 5 Click **Save**.

Step 6 Click **MAC Address Pool** and then **Add MAC Address Pool**, and configure the following fields.

For clustering in Individual interface mode, you can configure a MAC address pool for the interface. It is not common to manually configure MAC addresses for an interface, but if you have special needs to do so, then this pool is used to assign a unique MAC address to each interface. See Configure the MAC address, on page 1110.

- **Name**—Enter the name of the address pool. It can be up to 64 characters
- **Description**—Add an optional description for this pool.
- **MAC Address Range**—Enter a range of MAC addresses available in the pool. Use a dash between the beginning and the end address, for example: 000C.F142.4CD1-000C.F142.4CD7.
- **Allow Overrides**—Check this check box to enable overrides. Click the expand arrow to show the **Overrides** table. You can add a new override by clicking **Add**. See Object overrides, on page 2939 for more information.

Create an application filter

An application filter is a reusable object that allows you to organize and manage applications efficiently for use in policies that require application matching.

An application filter does these:

- Simplifies application control and policy management.
- Groups applications based on type, risk, business relevance, category, and tags.
- Supports policies such as access control, QoS, intelligent application bypass, and decryption.

Application filter configuration options

- Use the system-provided filters to find applications. Select the desired applications and click **Add to Rule**.
- To remove an application from the filter, click the delete icon next to it.
- Click the info icon next to an application to view its information.

For detailed information on application characteristics, refer to Application rule conditions, on page 1701.

Configure an AS PATH object

Enable control over BGP routing updates and filtering by specifying a sequence of autonomous system (AS) numbers.

An AS Path is a mandatory attribute to set up BGP. It is a sequence of AS numbers through which a network can be accessed. AS path is a sequence of intermediate AS numbers between source and destination routers that form a directed route for packets to travel. Neighboring autonomous systems (ASes) use BGP to exchange and update messages about how to reach different AS prefixes. After each router makes a new local decision on the best route to a destination, it sends the route and associated path information, including distance metrics and path attributes, to its peers. As this information travels through the network, each router along the path prepends its unique AS number to a list of ASes in the BGP message. This list is the route's AS-PATH.

An AS path, along with an AS prefix, provides a specific handle for a one-way transit route through the network. Use the Configure AS Path page to create, copy, and edit autonomous system (AS) path policy objects. Create AS path objects to use when configuring route maps, policy maps, or BGP Neighbor Filtering. An AS path filter allows you to filter the routing update message by using regular expressions.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Choose **Objects** > **Object Management** > **AS Path** from the table of contents.

Step 2 Click **Add AS PATH**.

Step 3 Enter a name for the AS PATH object in the **Name** field. Valid values are between 1 and 500.

Step 4 Click **Add** on the **New AS PATH Object** window.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.
b) Specify the regular expression that defines the AS PATH filter in the **Regular Expression** field.
c) Click **Add**.

Step 5 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 6 Click **Save**.

Configure a BFD template

Configure a BFD template to define BFD interval values and authentication settings for single-hop or multi-hop sessions. This enables flexible and reusable BFD configurations across interfaces.

The BFD template specifies a set of BFD interval values. BFD interval values as configured in the BFD template are not specific to a single interface. You can also configure authentication for single-hop and multi-hop sessions. Echo mode is disabled by default. You can enable Echo mode on single-hop only.

Before you begin

No prerequisites are required to configure a BFD template.

Procedure

Step 1 In the Cloud-Delivered Firewall Management Center, choose **Objects** > **Object Management** > **BFD Template**.

Step 2 Click **Add BFD Template** or **Edit**.

Note
If you are editing a template, you cannot modify its name and type.

Step 3 On the **Template** tab, configure the following:

- **Template Name**—The name of this BFD template. You must assign a name in order to configure the rest of the parameters in the template. The template name cannot have spaces and cannot have only numbers.
- **Type**—Click the **Single-Hop** or **Multi-Hop** radio button.
- **Enable Echo**—(Optional) Enables Echo for the single-hop template.

If the Echo function is not negotiated, BFD control packets are sent at a high rate to meet the detection time. If the Echo function is negotiated, BFD control packets are sent at a slower, negotiated rate and self-directed echo packets are sent at a high rate. We recommend that you use Echo mode, if possible.

Step 4 On the **Interval** tab, configure the following:

a) From the **Interval Type** drop-down list, select **Microseconds** or **Milliseconds**.
b) In the **Multiplier** field, enter the value to be used for computing the hold down time. This value indicates the number of consecutive BFD control packets that must be missed from a BFD peer before BFD declares that the peer is unavailable and the Layer 3 BFD peer is informed of the failure. The range is 3 to 50. The default is 3.
c) In the **Minimum Transmit** field, enter the minimum transmit interval capability. The range is 50 to 999 milliseconds or 50000 to 999000 microseconds.
d) In the **Minimum Receive** field, enter the minimum receive interval capability. The range is 50 to 999 milliseconds or 50000 to 999000 microseconds.

Step 5 On the **Authentication** tab, configure the following:

- **Authentication Type**—Select **NONE**, **md5**, **meticulous-sha-1**, **meticulous-md5**, or **sha-1** from the drop-down list.
- **Encrypted Password**—(Optional) Enables encryption of the authentication password.
- **Password**—The authentication password that must be sent and received in the packets using the routing protocol being authenticated. The valid value is a string containing 1 to 29 uppercase and lowercase alphanumeric characters, except that the first character cannot be a digit or a digit followed by a whitespace. For example, '1password' or '0 password' is invalid.
- **Key ID**—The shared key ID that matches the key value. The range is 0 to 255.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the BFD template configuration.

The BFD template is configured with the specified interval values, authentication settings, and echo mode options. The template can now be applied to interfaces as needed.

Cipher suite lists

A cipher suite list is an object that

- contains several cipher suites,
- enables negotiation of SSL or TLS encrypted sessions, and
- can be used in SSL rules to control encrypted traffic based on the negotiated cipher suite.

Usage of cipher suite lists in SSL rules

You can use cipher suites and cipher suite lists in SSL rules to control encrypted traffic. These rules evaluate whether the client and server negotiated the SSL session using that cipher suite. If you add a cipher suite list to an SSL rule, and if an SSL session negotiates with any cipher suite from the list, the rule matches the session.

Note You can use cipher suites in the web interface wherever you use cipher suite lists. However, you cannot add, change, or remove cipher suites.

Create a cipher suite list

Create a custom cipher suite list to control which cipher suites are available for use in your environment.

Use this task when defining or updating the set of cipher suites to align with security policies or updated encryption standards.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Cipher Suite List** from the list of object types.

Step 2 Click **Add Cipher Suites**.

Step 3 Enter a **Name**.

Step 4 Choose one or more cipher suites from the **Available Ciphers** list.

Step 5 Click **Add**.

Step 6 Optionally, click **Delete** ( ) next to any cipher suites in the **Selected Ciphers** list that you want to remove.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Configure a community list

Configure a community list to group BGP destinations that share common attributes for route tagging and policy enforcement. This task enables you to create, copy, and edit community list policy objects for use in route maps or policy maps.

Community lists allow marking sets of prefixes for common routing policies, such as filtering or assigning local preference.

A community is an optional transitive BGP attribute used for route tagging. Communities are numerical values assigned to prefixes. They are advertised to neighbors. Upstream providers use these markers to apply routing policies, such as filtering or modifying attributes.

Use the Configure Community Lists page to create, copy, and edit community list policy objects. You can use this object with Firewall Threat Defense devices.

Community lists are used in match clauses of route maps and consist of ordered matching statements. Each destination is compared against the rules. The system stops processing further rules when it finds a match.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Community List** from the table of contents.

Step 2 Click **Add Community List**.

Step 3 In the **Name** field, specify a name for the community list object.

Step 4 Click **Add** on the **New Community List Object** window.

Step 5 Select the **Standard** radio button to indicate the community rule type.

Standard community lists are used to specify well-known communities and community numbers.

Note
You cannot have entries using Standard and entries using Expanded community rule types in the same Community List object.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.
b) In the **COMMUNITIES** field, specify a community number. Valid values can be from 1 to 4294967295 or from 0:1 to 65534:65535.
c) Select the appropriate **Route Type**.

- Internet—Select to specify the Internet well-known community. Routes with this community are advertised to all peers (internal and external).
- No Advertise—Select to specify the no-advertise well-known community. Routes with this community are not advertised to any peer (internal or external).
- No Export—Select to specify the no-export well-known community. Routes with this community are advertised to only peers in the same autonomous system or to only other sub-autonomous systems within a confederation. These routes are not advertised to external peers.

Step 6 Select the **Expanded** radio button to indicate the community rule type.

Expanded community lists filter communities with regular expressions, which specify patterns used to match community attributes.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.
b) Specify the regular expression in the **Expressions** field.

Step 7 Click **Add**.

Step 8 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 9 Click **Save**.

Configure an extended community list

Create an extended community list to organize destinations by shared attributes, enabling precise BGP route filtering and policy object use.

An extended community is a larger group of destinations that share some common attribute. The BGP extended community list has attributes that can be used to mark a set of prefixes that share a common attribute. These markers are used in the match clause of a route map to filter the routes for implementing route leaks among virtual routers. You can also define policy list objects with the extended community list for filtering. Matching statements are organized in an ordered list within the extended community list. Routes are matched against the rules until a match is found with the specified route target (standard rule type) or regular expression (expanded rule type). Create and edit extended community list policy objects using the Extended Community page.

Note The extended community lists are applicable only for configuring import or export of routes.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** > **Community List** > **Extended Community**, and then click **Add Extended Community List**.

Step 2 In the **Name** field, specify a name for the extended community list object. You can enter up to 80 characters for the name.

Step 3 Select the extended community rule type:

- To specify one or more route targets, click the **Standard** radio button .
- To specify regular expressions, click the **Expanded** radio button.

Note
Use either Standard or Expanded extended community rule type in each Extended Community List object.

Step 4 Click **Add**.

Step 5 If you have selected **Standard** as the extended community rule type, specify the following:

a) In the **Sequence No** field, enter the order in which you want the rule to be executed.

The sequence number must be unique in the list.

b) From the **Action** drop-down list, if you want to permit routes that have matching route target that is specified here, select **Allow**; if you want to deny routes that have matching route target that is specified here, select **Block**.

c) In the **Route Target** field, specify a route target.

- You can add a single route target or a set of route targets separated by commas in a single entry. For example, *1:2,1:4,1:6*.
- Valid values can be from 1:1 to 65534:65535.
- You can have a maximum of 8 route targets in an entry.
- You cannot have redundant route target set across multiple entries. For example, say you want to configure *seq1* with *1:200,100:100,1:300* route targets, and *seq2* with *1:300,100:100,1:200* route targets. This results in redundant route target set and cannot be deployed.

Step 6 If you have selected **Expanded** as the extended community rule type, specify the following:

a) In the **Sequence No** field, enter the order in which you want the rule to be executed.

The sequence number must be unique in the list.

b) From the **Action** drop-down list, if you want to permit routes that have matching regular expression that is specified here, select **Allow**; if you want to deny routes that have matching regular expression that is specified here, select **Block**.

c) Specify the regular expression in the **Expressions** field.

- You can add a single route target or a set of route targets separated by a space in a single entry. For example, *^((16) | (18)):(.)$*.
- You can add a maximum of 16 regular expressions to an entry.
- You cannot have redundant regular expression set across multiple entries. For example, say you want to configure *seq1* with *^((16) | (18)):(.)$ ^4_[0-9]*$* route targets, and *seq2* with *^4_[0-9]*$ ^((16) | (18)):(.)$* route targets. This results in redundant regular expression set and cannot be deployed.

For details on BGP regular expressions, see here.

Step 7 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 8 Click **Save**.

The extended community list can be referenced in the match clause of the Route Map object or Policy List object:

- In the Route Map object, the name of the extended community list is displayed in the **Add Route Map Entry** > **Match Clause** > **BGP** > **Community List** > **Add Extended Community List** dialog. For more details on configuring BGP settings in a route map, see Configure a route map, on page 3017.
- In the Policy List object, the name of the extended community list is displayed in the **Add Policy List** > **Community Rule** > **Add Extended Community List** dialog. For more details on configuring BGP settings in a policy list, see Configure a policy list, on page 3012.

DHCP IPv6 pools

A DHCP IPv6 pool is a configuration element that

- defines a range of IPv6 addresses for DHCP assignment,
- is referenced when configuring DHCPv6 services, and
- enables address allocation to IPv6 clients.

For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature (Enable the IPv6 prefix delegation client, on page 1093), you can configure the Firewall Threat Defense to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the Firewall Threat Defense by defining a DHCP IPv6 Pool and assigning it to the DHCPv6 server. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses based on prefixes received in Router Advertisement messages; in other words, based on the prefix that the Firewall Threat Defense received using Prefix Delegation.

To add a pool, see Create a DHCP IPv6 pool, on page 1229.

Distinguished names

A distinguished name is a certificate attribute that

- represents the subject or issuer for a public key certificate,
- can be used in rules to control encrypted traffic based on certificate negotiation, and
- consists of attributes such as country code, common name, organization, and organizational unit.

A *distinguished name group* is a named collection of existing distinguished name objects.

Distinguished name attributes and usage

Distinguished name objects and distinguished name groups can be used in TLS/SSL rules to control encrypted traffic based on whether the client and server negotiated the TLS/SSL session using a server certificate with the distinguished name as subject or issuer.

A distinguished name can include a country code, common name, organization, and organizational unit. In most cases, the certificate includes only a common name. For example, the common name in the certificate for `https://www.cisco.com` is `cisco.com`. The certificate can include multiple Subject Alternative Names (SANs). You can use these names as distinguished names (DNs) in a rule condition. For detailed information about SANs, see RFC 5280, section 4.2.1.6.

The format of a distinguished name object that references a common name is `CN=`*`name`*. If you add a DN rule condition without `CN=`, the system prepends `CN=` before saving the object.

The system uses Server Name Indication (SNI) to match the DN in the TLS/SSL rule whenever possible.

You can also add a distinguished name with one of each of the attributes listed in this table, separated by commas.

Table 160: Distinguished name attributes

Attribute	Description	Allowed Values
C	Country Code	two alphabetic characters
CN	Common Name	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces
O	Organization	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces
OU	Organizational Unit	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces

DN rule conditions

- When the system first detects an encrypted session to a new server, DN data is not available for ClientHello processing. As a result, the first session might not be decrypted.

  If the server requests TLS 1.3, the setting for TLS server identity discovery can help by making sure the server certificate is known before making decryption policy decisions. For more information, see TLS server identity discovery, on page 2692.

- You *cannot* configure a distinguished name condition if you also choose the **Decrypt - Known Key** action. Because that action requires you to choose a server certificate to decrypt traffic, the certificate already matches the traffic.

Wildcard examples

You can define one or more asterisks (*) as wildcards in an attribute. In a common name attribute, you can define one or more asterisks per domain name label. Wildcards match only in that label, but you can define multiple labels with wildcards. See this table for examples.

Table 161: Common name attribute wildcard examples

Attribute	Matches	Does not match
CN=*ample.com	example.com	mail.example.com example.text.com ampleexam.com
CN=exam*.com	example.com	mail.example.com example.text.com ampleexam.com
CN=*xamp*.com	example.com	mail.example.com example.text.com ampleexam.com

Attribute	Matches	Does not match
CN=*.example.com	mail.example.com	www.myhost.example.com example.com example.text.com ampleexam.com

Note The DN object `CN=amp.cisco.com` would *not* match a CN like `CN=auth.amp.cisco.com`, which is why wildcards are recommended in these cases.

For more information and examples, see Conditions for distinguished name rules in decryption policies, on page 2208.

Related Topics

Conditions for distinguished name rules in decryption policies, on page 2208

Create a distinguished name object

Enables you to create distinguished name objects, which uniquely identify and reference network entities within Secure Firewall Management Center.

Distinguished name objects are used to specify unique identifiers for certificates or access policies. Perform this task when configuring objects that require a distinguished or common name for identification.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Distinguished Name** node, and choose **Individual Objects**.

Step 3 Click **Add Distinguished Name**.

Step 4 Enter a **Name**.

Step 5 In the **DN** field, enter a value for the distinguished name or common name. You have the following options:

- If you add a distinguished name, you can include one of each attribute listed in Distinguished names, on page 2962 separated by commas.
- If you add a common name, you can include multiple labels and wild cards.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

DNS server groups

A DNS server group is a network configuration object that

- contains multiple Domain Name System (DNS) servers
- resolves fully qualified domain names (FQDNs), such as `www.example.com`, to IP addresses, and
- enables efficient domain name resolution within a network.

Create a DNS server group object

Create and manage a group of DNS servers for use in network policies to simplify DNS server assignments.

DNS server group objects allow you to define multiple DNS servers that can be referenced by other configuration elements. Use this task to centralize DNS server assignments for interfaces or policies.

Procedure

Step 1 Choose **Objects** > **Object Management** > **DNS Server Group** from the network objects list.

Step 2 Click **Add DNS Server Group**.

Step 3 Enter a **Name**.

Step 4 Optionally, enter the **Default Domain** that will be used to append to the host names that are not fully-qualified.

This setting is only used for the default server group.

Step 5 The default **Timeout** and **Retries** values are pre-populated. Change these values if necessary.

- Retries—Specify how many times (0 to 10) to retry contacting DNS servers if there is no response. Default is 2.
- Timeout—Specify how many seconds (1 to 30) to wait before trying the next DNS server. Default is 2 seconds; this timeout doubles with each retry.

Step 6 Enter the **DNS Servers** IP addresses (IPv4 or IPv6), separated by commas. *(Maximum of 6 servers per group.)*

Step 7 Click **Save**.

What to do next

Assign the configured DNS server group to interface objects in the DNS platform settings. For more information, see Configure DNS server settings, on page 1246.

External attributes

Dynamic objects

A dynamic object is a network object that

- specifies one or many IP addresses retrieved using REST API calls or the Dynamic Attributes Connector from cloud sources,
- can be used in access control or DNS rules without deploying dynamic changes, and
- automatically updates object values on managed devices after being pushed by the Dynamic Attributes Connector.

Kinds of dynamic objects

There are these kinds of dynamic objects:

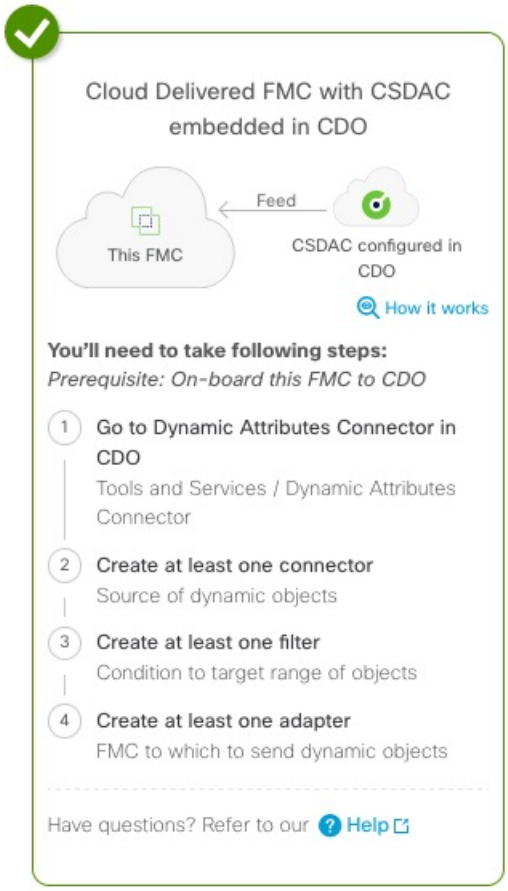
- Connector-created dynamic objects: The system pushes dynamic objects created using the dynamic attributes connector to the Cloud-Delivered Firewall Management Center as soon as you create them. The system updates them at regular intervals.
- API-created dynamic objects have the following characteristics:
  - Are IP addresses, with or without classless inter-domain routing (CIDR), that can be used in access control rules much like a network object.
  - Do not support fully-qualified domain names or address ranges.
  - You must update these objects using an API.

For more information about API-created dynamic objects, see API-created dynamic objects, on page 2971.

Note Unlike most other objects, dynamic objects do not have to be deployed to managed devices to take effect. Add a dynamic object to the **Dynamic Attributes** tab in a rule and deploy the rule. The object values are automatically updated on the managed device as soon as possible after being pushed by the Dynamic Attributes Connector.

Create dynamic objects with Cloud-Delivered Firewall Management Center

This page appears if you indicated that you are configuring the Dynamic Attributes Connector with Cloud-Delivered Firewall Management Center in Security Cloud Control.



A dynamic object is a firewall configuration element that

- enables automatic retrieval of IP addresses from cloud services,
- allows for dynamic updates within access control rules, and
- provides integration with external systems for adaptive security policies.

To deploy this feature:

1. Configure connectors: Retrieve IP addresses from supported cloud services for automatic policy updates.

   For more details, see Connectors, on page 2631.
2. Configure dynamic attributes filters: Specify which IP addresses should be sent to the management center.

   For more details, see Create dynamic attributes filters, on page 2455.
3. Configure a Cloud-Delivered Firewall Management Center adapter: Ensure that the retrieved IP addresses are sent to the correct device.
4. Access **Administration** > **Firewall Management Center**: Use the Tools menu to select **Cloud-Delivered FMC**.
5. View your dynamic objects: Navigate to **Objects** > **Object Management** > **External Attributes** > **Dynamic Object** to see the list.

6. Use dynamic objects in access control rules: In access control policies, use the **Dynamic Attributes** tab to add dynamic objects.

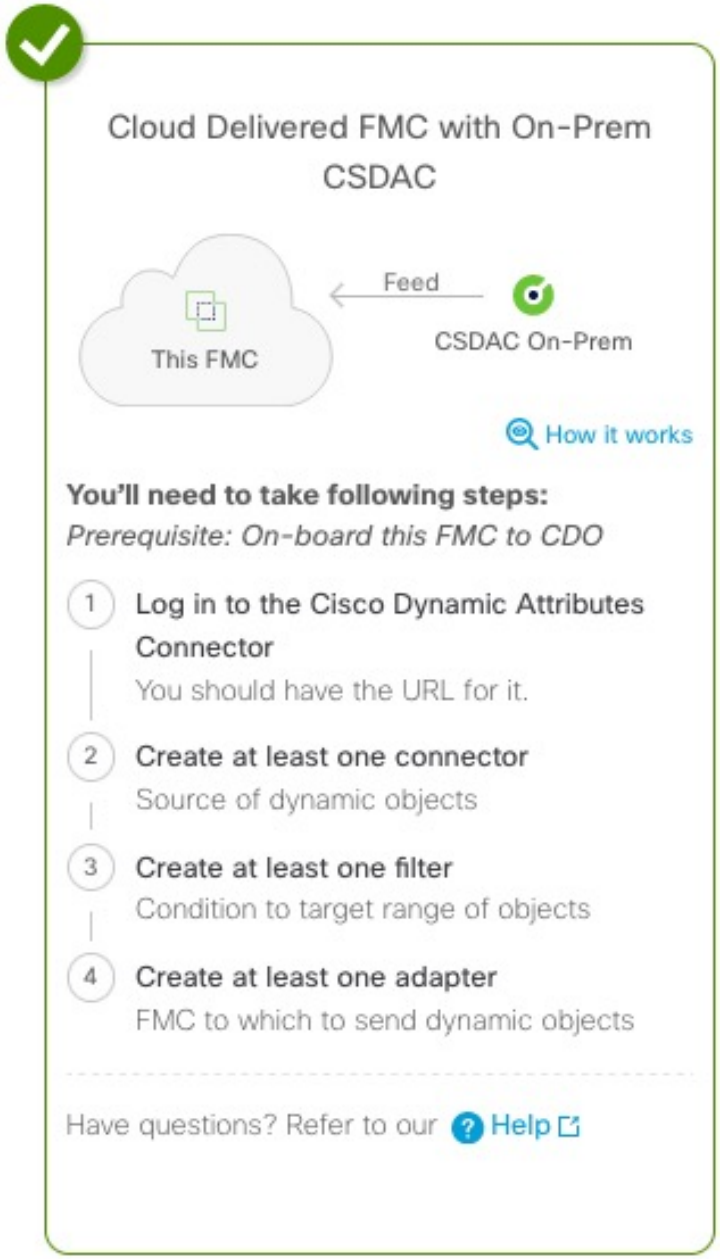
   You can also use the objects in DNS rules.

Create Dynamic Objects with Cloud-Delivered Firewall Management Center and On-Premises Dynamic Attributes Connector

Dynamic objects are automated network objects that synchronize IP addresses from cloud services to your management center via the dynamic attributes connector.

This page appears when you choose to configure an on-premises Dynamic Attributes Connector with Cloud-Delivered Firewall Management Center in Security Cloud Control.

Figure 397: Cloud-Delivered FMC and Dynamic Attributes Connector Integration



1. Install the Dynamic Attributes Connector on a supported Linux virtual machine.
2. Configure connectors. These connectors retrieve IP addresses from cloud services.
3. Configure adapters. These adapters send IP addresses to your Cloud-Delivered Firewall Management Center.
4. Configure dynamic attributes filters. These filters determine what IP addresses to send to the management center.

5. Click **Administration** > **Firewall Management Center** and choose **Cloud-Delivered FMC**.
6. View your dynamic objects at **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.
7. Use dynamic objects in access control rules (**Policies** > **Access Control heading** > **Access Control**, then click the **Dynamic Attributes** tab).

   You can also use the objects in DNS rules.

Dynamic objects page features

The Dynamic Objects page displays all dynamic objects created using the API or the Cisco Secure Dynamic Attributes Connector. You can edit, delete, or view IP address mappings on this page. The page enables you to view or download IP addresses associated with each object.

The page available at **Objects** > **Object Management** > **External Attributes** > **Dynamic Object** is displayed similar to this if you have already configured some dynamic objects.



This page displays information about each dynamic object and enables you to view or download IP addresses associated with that object. For more information, see Dynamic object mappings, on page 2970.

Dynamic object mappings

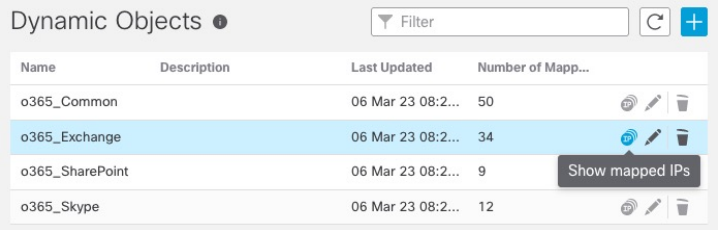
A dynamic object mapping is a configuration feature that

- enables connectors to send IPs matching dynamic attribute filters to the Cloud-Delivered Firewall Management Center,
- allows users to view or download a current list of mapped IP addresses, and
- updates IP addresses dynamically over time to reflect changes in access control rules.

Mapped IP address information

If you configured dynamic objects through the API or by using the dynamic attributes connector, the connectors send IPs that match dynamic attribute filters to the Cloud-Delivered Firewall Management Center at regular intervals.

To view or download a current list of these IP addresses, click **Show Mapped IDs** in the Cloud-Delivered Firewall Management Center, as illustrated.



IP addresses are added dynamically over time. Review the mapped IP addresses regularly, especially if the access control rules do not behave as expected.

Related topics

- API-created dynamic objects, on page 2971
- About the Dynamic Attributes Connector , on page 2621

API-created dynamic objects

An API-created dynamic object is a network object that

- specifies one or more IP addresses retrieved using REST API calls or the Dynamic Attributes Connector,
- can be used in access control or DNS rules without deploying dynamic changes to the objects, and
- updates IP addresses from cloud sources.

For more information about the dynamic attributes connector, see the information later in this guide.

Differences between dynamic objects and network objects

- Dynamic objects created using the dynamic attributes connector are pushed to the Cloud-Delivered Firewall Management Center as soon as they are created. These objects are then updated at regular intervals.
- API-created dynamic objects:
  - are IP addresses, with or without classless inter-domain routing (CIDR), that can be used in access control rules much like a network object,
  - do not support fully-qualified domain names or address ranges, and
  - must be updated using an API.

Add or edit an API-created dynamic object

This task enables you to add or edit a dynamic object using the API, allowing you to group IP addresses for use in access control rules. Dynamic objects help automate address management and streamline policy configuration.

A dynamic object is a group of IP addresses, with or without classless inter-domain routing (CIDR). You can use dynamic objects in access control rules much like a network object. This procedure is relevant when you need to manage dynamic address groups via the API instead of manually creating static objects.

Note This procedure is not necessary if you use the Dynamic Attributes Connector because it automatically creates dynamic objects for you.

Before you begin

Consult the *Firepower Management Center REST API Quick Start Guide* for information about how to use the object services REST API to populate the IP object with an address. Dynamic objects do not require deployment.

Follow these steps to add or edit an API-created dynamic object:

Procedure

Step 1 Click **Objects** > **Object Management**.

Step 2 Click **External Attributes** > **Dynamic Objects**.

Step 3 Click **Add Dynamic Object** or **Edit** ( ).

Step 4 Enter a **Name** for the object and an optional **Description**.

Step 5 From the **Type** list, click **IP**.

What to do next

If necessary, update the dynamic object using the API. Deployment is not required.

Security group tags

A security group tag is a network security object that

- specifies a single SGT value,
- is used in rules to control traffic with SGT attributes not assigned by Cisco ISE, and
- cannot be grouped or overridden.

Related Topics

Autotransition from custom SGTs to ISE SGTs
Custom SGT conditions
ISE SGTs versus custom SGT rule conditions

Create a security group tag object

This task allows you to add a security group tag object to enforce network access policies based on SGT.

You can create security group tag objects only in the global domain. A Control license is required for Classic devices. Any Smart License provides coverage for Smart Licensed devices.

Before you begin

- Disable ISE or ISE-PIC connections because custom SGT objects cannot be created if either is used as an identity source.

  Click **Integration** > **Other Integrations** > **Identity Sources**, click **None**, then click **Save**.

- See Best practices for ISE/ISE-PIC integration, on page 2376 for guidelines for using the identity source.

Follow these steps to create a security group tag object:

Procedure

Step 1 Click **Objects** > **Object Management**.
Step 2 Click **External Attributes** > **Security Group Tag**.
Step 3 Click **Add Security Group Tag**.
Step 4 Enter a **Name**.
Step 5 Optionally, enter a **Description**.
Step 6 In the **Tag** field, enter a single SGT.
Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

File lists

A file list is a threat management feature that

- allows you to manually specify files using unique SHA-256 hash values,
- overrides AMP cloud verdicts to improve detection accuracy within your environment, and
- supports categorization into clean or custom detection lists, enabling precise control over file treatment.

If you use malware defense and the AMP cloud incorrectly identifies a file's disposition, you can add the file to a file list to improve future detection. Specify these files using SHA-256 hash values. You can add up to 10,000 unique SHA-256 values in each file list.

File list categories

File lists are organized into two predefined categories.

- **Clean list**: If you add a file to this list, the system treats it as if the AMP cloud assigned a clean disposition.
- **Custom detection list**: If you add a file to this list, the system treats it as if the AMP cloud assigned a malware disposition.

Because you manually specify the blocking behavior for the files included in these lists, the system does not query the AMP cloud for these files' dispositions. You must configure a rule in the file policy with either the **Malware Cloud Lookup** or **Block Malware** action and a matching file type to calculate a file's SHA value.



Caution Do **not** include malware on the clean list. The clean list overrides both the AMP cloud disposition and the custom detection list.

Source files for file lists

Source files allow you to add multiple SHA-256 values to a file list by uploading a comma-separated value (CSV) file. The system validates the contents and populates the file list with valid SHA-256 values. Upload source files must meet the following requirements:

- The source file must be a simple text file with a .CSV extension.
- Any header must start with a pound sign; it is treated as a comment and will not be uploaded.
- Each entry should contain a single SHA-256 value followed by a description, and must end with either LF or CR+LF newline character.
- The system ignores any additional information present in the entry.

Note the following behaviors and limitations:

- Deleting a source file from the file list also removes all associated SHA-256 hashes from the list.
- You cannot upload multiple files to a file list if the upload would result in more than 10,000 distinct SHA-256 values in the list.
- The system truncates descriptions exceeding 256 characters to the first 256 characters when uploading. If the description contains commas, you must use an escape character (`\,`). If no description is provided, the source file name is used instead.
- All non-duplicate SHA-256 values are added to the file list. If a SHA-256 value already exists, it is not modified by subsequent uploads. Any threat name or description is derived from the individual SHA-256 value in related events.
- Invalid SHA-256 values in a source file are not uploaded.
- If multiple uploaded source files contain entries for the same SHA-256 value, the system uses the most recent value.
- If a source file contains multiple entries for the same SHA-256 value, the system uses the last entry.
- You cannot edit a source file directly within the object manager. To make changes, edit your local source file, delete the existing file from the system, then upload the revised file.
- The number of entries associated with a source file refers to the number of distinct SHA-256 values. Deleting a source file decreases the file list's SHA-256 entries by the number of valid entries in the deleted file.

Add individual SHA-256 values to file lists

Add unique SHA-256 values to file lists to enhance file detection and management within the system.

You must have the Malware Defense license for this procedure.

Adding SHA-256 values identifies files for detection and prevents duplicates by ensuring only unique values are included.

Before you begin

- Right-click a file or malware event from the event view, choose **Show Full Text** in the context menu, and copy the full SHA-256 value for pasting into the file list.

Follow these steps to add individual SHA-256 values to file lists:

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** ( ) next to the clean list or custom detection list where you want to add a file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Choose `Enter SHA Value` from the **Add by** drop-down list.

Step 5 Enter a description of the source file in the **Description** field.

Step 6 Enter or paste the file's entire value in the **SHA-256** field. The system does not support matching partial values.

Step 7 Click **Add**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Note After configuration changes are deployed, the system no longer queries the AMP cloud for files on the list.

Upload an individual file to a file list

You must have the Malware Defense license for this procedure.

If you have a copy of the file you want to add to a file list, you can upload the file to the Cloud-Delivered Firewall Management Center for analysis. The system automatically calculates the file's SHA-256 value on upload. No file size limit applies for SHA-256 calculation.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** ( ) next to the clean list or custom detection list where you want to add a file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 From the **Add by** drop-down list, choose **Calculate SHA**.

Step 5 (Optional) Enter a description of the file in the **Description** field. If you leave this field blank, the file name becomes the description during upload.

Step 6 Click **Browse**, and choose a file to upload.

Step 7 Click **Calculate and Add SHA** to upload the file and calculate the SHA-256 value.

Step 8 Click **Save**.

After completing these steps, the file is uploaded, its SHA-256 value is calculated, and the file is added to the selected file list for analysis.

What to do next

- If an active policy references your object, deploy configuration changes.

Note After you deploy configuration changes, the system no longer queries the AMP cloud for files on the list.

Upload source files to file lists

This task enables you to upload source files to file lists, allowing for efficient management and use of file data within the system.

You must have the Malware Defense license for this procedure.

Use this task when you need to add values from a source file to an existing file list in the system.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **File List**.

Step 3 Click **Edit** ( ) next to the file list where you want to add values from a source file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 In the **Add by** drop-down list, choose `List of SHAs`.

Step 5 (Optional) Enter a description of the source file in the **Description** field. If you do not enter a description, the system uses the file name.

Step 6 Click **Browse** to browse to the source file, then click **Upload and Add List**.

Step 7 Click **Save**.

After you complete these steps, the system uploads the source file and adds its values to the selected file list.

What to do next

- If an active policy references your object, deploy configuration changes.



Note After you deploy the policies, the system stops querying the AMP cloud for files on the list.

Edit SHA-256 values in file lists

You must have the Malware Defense license for this procedure.

You can edit or delete individual SHA-256 values on a file list. However, you cannot directly edit a source file within the object manager. To make changes, first modify your source file, delete the copy on the system, and then upload the modified file.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **File List**.

Step 3 Click **Edit** () next to the clean list or custom detection list where you want to modify a file.

If **View** () appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 You can:

- Click **Edit** () next to the SHA-256 value you want to change, and modify the **SHA-256** or **Description** values as desired.
- Click **Delete** () next to the SHA-256 value you want to delete.

Step 5 Click **Save** to update the file entry in the list.

Step 6 Click **Save** to save the file list.

What to do next

- If an active policy references your object, deploy configuration changes.

Note After configuration changes are deployed, the system no longer queries the AMP cloud for files on the list.

Download source files from a file list

This task allows you to download source files from file lists, enabling you to obtain files for analysis or operational needs.

You must have the Malware Defense license for this procedure.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** ( ) next to the clean list or custom detection list where you want to download a source file.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Next to the source file you want to download, click **View** ( ).

Step 5 Click **Download SHA List** and follow the prompts to save the source file.

Step 6 Click **Close**.

FlexConfig objects

A FlexConfig object is a configuration element that

- provides customized configuration of features on Firewall Threat Defense devices
- enables configuration of features not otherwise available through Cloud-Delivered Firewall Management Center, and
- uses policy objects in FlexConfig policies to customize device behavior.

FlexConfig object types

You can configure these types of objects for FlexConfig.

- **Text objects** define free-form text strings used as variables in a FlexConfig object. These objects can have a single value or multiple values. Several predefined text objects are included in FlexConfig objects. To customize device behavior, edit the contents of the text object when using a FlexConfig object. When editing a predefined object, create device overrides for each device instead of changing the default values. This helps prevent configuration issues for other users who use the same FlexConfig object. For information on configuring text objects, see Configure FlexConfig text objects, on page 1337.

- **FlexConfig objects** include device configuration commands, variables, and scripting language instructions. During deployment, the system processes these instructions to create configuration commands with customized parameters for the target devices. You can add instructions at the beginning (prepended) of the configuration or at the end (appended), based on your requirements. Append any FlexConfig that depends on Cloud-Delivered Firewall Management Center-configured objects (for example, a network object) to the configuration deployment. This ensures the necessary objects are configured before the FlexConfig refers to them. For more information on configuring FlexConfig objects, see Configure FlexConfig objects, on page 1333.

Geolocation objects

A geolocation object is a network policy element that

- represents one or more countries or continents identified as the source or destination of traffic on your monitored network,
- can be used in various places in the system's web interface, including access control policies, SSL policies, remote access VPN, and event searches, and
- enables you to write access control rules that block traffic to or from certain countries.

Geolocation database updates

To ensure that you are using up-to-date information to filter your network traffic, we strongly recommend that you regularly update your Geolocation Database (GeoDB).

Create a geolocation object

Create a geolocation object to define groups of countries and continents for use in security policies.

Use geolocation objects in access control and policy rules to manage network traffic based on geographic criteria.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Geolocation** from the list of object types.

Step 3 Click **Add Geolocation**.

Step 4 Enter a **Name**.

Step 5 Select the countries and/or continents to include.

- Selecting a continent includes all countries within that continent, including future additions via GeoDB updates.
- You can select any combination of countries and continents.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Interfaces

An interface is a network connection point that

- can be assigned to a *security zone*, an *interface group*, or both,
- enables application of security policies based on zones or groups, and
- supports flexible traffic control between different network segments.

Security zone and interface group assignments

Each interface can be assigned to a *security zone*, an *interface group*, or both. Apply your security policy based on zones or groups. For example, assign the "inside" interface to the "inside" zone and the "outside" interface to the "outside" zone. Configure the access control policy to allow traffic from inside to outside and block traffic from outside to inside. Some policies support only security zones. Other policies support both zones and groups.

For more information about interface objects, see Security zones and interface groups, on page 999.

To add interface objects, see Create security zone and interface group objects, on page 1002.

Key chains

Key chains help you rotate authentication keys for IGP peers to enhance data security, prevent unauthorized users from guessing routing protocol authentication keys, and maintain uninterrupted, secure communication by allowing overlapping key lifetimes.

Key chain authentication and key management

Use rotating keys in key chains, with a duration of 180 days or less, to protect your devices and data. Rotating keys apply only to the OSPFv2 protocol.

Note Only the MD5 cryptographic algorithm is used for authentication.

Each key in a key chain has two lifetimes:

- Accept lifetime: The time interval within which the device accepts the key during key exchange with another device.
- Send lifetime: The time interval within which the device sends the key during key exchange with another device.

During a key's send lifetime, the device sends routing update packets using that key. If the key is not within the accept lifetime, the device does not accept communication from other devices.

If lifetimes are not configured, the system treats the configuration as using an MD5 authentication key without defined timelines.

When configuring authentication for routing protocols that provide key chains, configure the keys in your key chain with overlapping lifetimes to prevent losing secure communication if no active key is available.

Key selection in a key chain is determined as follows:

- When a key chain has more than one valid key, OSPF selects the key with the longest lifetime.
- The system prefers a key with an infinite lifetime.
- If keys have the same lifetime, then key with the higher key ID is preferred.

Create a key chain object

Create key chain objects that securely manage and organize authentication keys for devices on your network.

Use this task when you need to define and manage key chains for secure authentication. Key chains enable secure key exchange, helping prevent unauthorized access and ensuring reliable communication between devices.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Key Chain** from the list of object types, and then click **Add Key Chain**.

Step 3 In the Add Key Chain Object dialog box, do the following:

a. Enter a name for the key chain in the **Name** field.

The name must start with an underscore or letter, followed by alphanumeric characters or one of the following special characters: hyphen, underscore, plus sign, or period.

b. Add a key to the key chain, click **Add**.

c. Specify the key identifier (between 0 and 255; use 0 only to signal an invalid key) in the **Key ID** field.

d. The **Algorithm** field and the **Crypto Encryption Type** field displays the supported algorithm and the encryption type, namely MD5 and Plain Text.

e. Enter the password in the **Crypto Key String** field and confirm it in the appropriate field. Passwords can be up to 80 characters and must not be a single digit or start with a digit followed by a space (e.g., "0 pass" or "1" are invalid).

f. Set the **Accept Lifetime** and **Send Lifetime** values to define the time intervals for key exchange. The times default to UTC. When specifying start and end lifetimes, ensure that the start is earlier than the end and is never left blank if the end is entered.

g. Click **Add** to create keys. Create at least two keys with overlapping lifetimes to prevent loss of secure communication.

Repeat steps 3b to 3g to create keys.

Step 4 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.
- If you want to add override values to this object, expand the Override section and click **Add**; see Add object overrides, on page 2942.

Step 5 Click **Save**.

After you save, you can use the key chain object in device authentication and secure communication configurations.

What to do next

- If an active policy references your object, deploy configuration changes.

Networks

A network object is a network entity that

- represents one or more IP addresses,
- can be used in access control policies, network variables, identity rules, network discovery rules, event searches, and reports, and
- is filtered automatically to show only valid objects for each configuration option.

Network object types

You can define a network object as one of these types:

- **Host**: A single IP address.

  IPv4 example:

  `209.165.200.225`

  IPv6 example:

  `2001:DB8::0DB8:800:200C:417A` or `2001:DB8:0:0:0DB8:800:200C:417A`

- **Range**: A range of IP addresses.

  IPv4 example:

  `209.165.200.225-209.165.200.250`

  IPv6 example:

  `2001:db8:0:cd30::1-2001:db8:0:cd30::1000`

- **Network**: An address block, also known as a subnet.

  IPv4 example:

  `209.165.200.224/27`

IPv6 example:

```
2001:DB8:0:CD30::/60
```

Note Security Intelligence ignores IP address blocks using a `/0` netmask.

- **FQDN**: A single fully-qualified domain name (FQDN). You can limit FQDN resolution to IPv4 address only, IPv6 address only, or both types. FQDNs must begin and end with a digit or letter. They can contain only letters, digits, and hyphens internally.

  Example:

  ```
  www.example.com
  ```

Note You can use FQDN objects in access control rules and prefilter rules, or manual NAT rules, only. The rules match the IP address obtained for the FQDN through a DNS lookup. To use an FQDN network object, ensure you have configured the DNS server settings in DNS server groups, on page 2965 and the DNS platform settings in Configure DNS server settings, on page 1246.

You *cannot* use FDQN network objects in identity rules.

- **Group**: A group of network objects or other network object groups. You can create nested groups by adding a network object group to another. Nesting is supported up to 10 levels.

Note You can add up to 100 network literals in a network object. Each nested network object group contains up to 100 network literals.

If you are using Cloud-Delivered Firewall Management Center

When you create a network object or group, it is replicated in the **Objects > FTD Network Objects** page in Security Cloud Control and vice-versa.

You can use the objects when specifying networks while configuring other Security Cloud Control-managed products, such as ASA or FDM.

Changes made to network objects or groups in either list are reflected in both. Deleting an object or group from either list also deletes its corresponding object or group from the other list.

Exception: If an object created on the Security Cloud Control list has the same name as an existing object on the Cloud-Delivered Firewall Management Center list, the object will not be replicated on the Cloud-Delivered Firewall Management Center list.

Network wildcard masks

A network wildcard mask is an IP address mask that

- uses a discontinuous mask of bits to define network objects,
- enables you to create network objects with expanded subnet IP addresses, and
- appears as **Network Wildcard** in the **Type** column of the network object list.

Network wildcard mask usage and support

You can create and manage wildcard mask objects from the Network Wildcard Mask page in **Objects** > **Object Management** .

Standard network objects use contiguous masks, while wildcard network objects use discontinuous masks.

This table shows examples of IP addresses and indicates whether they are considered network wildcard objects:

Table 162: Network wildcard mask examples

Example IP address	Network wildcard?	Object type
192.0.0.0/8	No	Network
10.10.0.0/255.255.0.0	No	Network
10.10.0.10/255.255.0.255	Yes	Network Wildcard
72.0.240.10/255.255.240.255	Yes	Network Wildcard

Note Network wildcard objects and groups containing network wildcard objects are allowed only when configuring these policies:

- Prefilter policy
- Access control policy
- NAT policy

Guidelines and limitations

- To create network wildcard objects, in the Cloud-Delivered Firewall Management Center UI, choose **Objects** > **Object Management** > **Network** and click **Add Network** and then **Add Object**. Select the **Network** option and enter the value as expanded subnet mask fotmat, such as 10.0.10.10 or 255.255.0.255.
- Features supported: Object override, group object support, group object override, wildcard literals, and wildcard object import.
- Network wildcard objects are supported only for IPv4 addresses.
- Network wildcard objects are supported from Cloud-Delivered Firewall Management Center and Firewall Threat Defense versions 7.1 onwards.
- Network wildcard objects are supported only for Snort-3.

Create a network object

Create network objects to define network addresses, address ranges, or groups that can be used in security policies and access rules.

Network objects allow you to efficiently manage and reuse network definitions across your security policies. This task is relevant when you need to add, clone, or modify network objects in your management center.

Procedure

Step 1 If you are accessing network objects from Security Cloud Control, choose **Objects**.

Step 2 Choose **Objects** > **Object Management**, and then Choose **Network** from the list of object types.

Step 3 Choose **Add Object** from the **Add Network** drop-down menu.

Step 4 Alternatively, you can clone an existing network object and edit the parameters to create a new network object. Click the **Clone** icon on the existing network object that you want to clone.

Step 5 Enter a **Name** and optionally, enter a **Description**.

Step 6 In the **Network** field, select the required option and enter an appropriate value; see Networks, on page 2982.

Note
You can add up to 100 network literals in a network object. Additionally, each nested network object group can contain a maximum of 100 network literals.

Step 7 (FQDN objects only) Select the DNS resolution from the **Lookup** drop-down menu to determine whether you want the IPv4, IPv6, or both IPv4 and IPv6 addresses associated with the FQDN.

Step 8 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.
- If you want to add override values to this object, expand the Override section and click **Add**; see Add object overrides, on page 2942.

Step 9 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Import network objects

For details on importing network objects, see Import objects, on page 2932.

Network objects and groups

A network object or group is a network management operation that

- modifies or removes network objects or groups from the Object Management page in Cloud-Delivered Firewall Management Center,

- automatically updates or deletes the corresponding network objects or groups replicated on the Objects page in Security Cloud Control, and
- ensures changes made to these objects on the Security Cloud Control objects page are reflected for the corresponding objects in Cloud-Delivered Firewall Management Center.



Caution When you edit or delete a network object or group from the Object Management page in Cloud-Delivered Firewall Management Center, this change also modifies or deletes the corresponding network object or group replicated on the Objects page in Security Cloud Control. Similarly, changes you make to these objects on the Security Cloud Control objects page are reflected for the corresponding objects in Cloud-Delivered Firewall Management Center.

PKI objects

PKI objects are public key certificates and paired private keys that

- represent the security components required to support your deployment
- consist of certificate authority certificates for internal and trusted CA objects
- consist of server certificates for internal and external certificate objects, and
- contain private keys paired with certificates for internal objects.

PKI object usage for applications and SSL rules

PKI objects support various security configurations across your network deployment.

You can use PKI objects for these applications:

- ISE/ISE-PIC identity source: Use trusted certificate authority objects and internal certificate objects to configure a connection to ISE/ISE-PIC as an identity source.
- Captive portal authentication: Use internal certificate objects to authenticate the identity of your captive portal device when connecting to users' web browsers.
- Secure LDAP or AD connections: Use trusted certificate authority objects to configure realms for secure connections to LDAP or AD servers.

If you use PKI objects in SSL rules, you can match traffic encrypted with:

- the certificate in an external certificate object
- a certificate either signed by the CA in a trusted CA object, or within the CA's chain of trust

If you use PKI objects in SSL rules, you can decrypt:

- outgoing traffic by re-signing the server certificate with an internal CA object
- incoming traffic using the known private key in an internal certificate object

You can manually input certificate and key information, upload a file containing that information, or in some cases, generate a new CA certificate and private key.

When you view a list of PKI objects in the object manager, the system displays the certificate's Subject distinguished name as the object value. Hover your pointer over the value to view the full certificate Subject distinguished name. To view other certificate details, edit the PKI object.



Note The Cloud-Delivered Firewall Management Center and managed devices encrypt all private keys stored in internal CA objects and internal certificate objects with a randomly generated key before saving them. If you upload private keys that are password protected, the appliance decrypts the key using the user-supplied password, then reencrypts it with the randomly generated key before saving it.

PKI objects for certificate enrollment

A certificate enrollment object is a PKI configuration item that

- contains the Certification Authority (CA) server information and enrollment parameters required for creating Certificate Signing Requests (CSRs)
- obtains identity certificates from the specified CA, and
- may include certificate revocation information.

Internal certificate authority objects

An internal certificate authority object is a security certificate object that

- represents the CA public key certificate of a CA your organization controls
- consists of the object name, CA certificate, and paired private key, and
- can be used in SSL rules to decrypt outgoing encrypted traffic by re-signing the server certificate with the internal CA.

Internal certificate authority object usage and management

You can create an internal certificate authority object in these ways:

- import an existing RSA-based or elliptic curve-based CA certificate and private key
- generate a new self-signed RSA-based CA certificate and private key
- generate an unsigned RSA-based CA certificate and private key. You must submit a certificate signing request (CSR) to another CA to sign the certificate before using the internal CA object.



Note If you reference an internal CA object in a **Decrypt - Resign** SSL rule and the rule matches an encrypted session, the user's browser may warn that the certificate is not trusted while negotiating the SSL handshake. To avoid this, add the internal CA object certificate to either the client or domain list of trusted root certificates.

After you create an internal CA object containing a signed certificate, you can download the CA certificate and private key. The system encrypts downloaded certificates and private keys with a user-provided password.

Whether system-generated or user-created, you can modify the internal CA object name, but cannot modify other object properties.

You cannot delete an internal CA object that is in use. Additionally, after you edit an internal CA object used in an SSL policy, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

CA certificates and private key import

A CA certificate and private key import is a security configuration process that

- enables uploading of X.509 v3 CA certificates and private keys,
- supports files encoded in DER or PEM formats, and
- validates that the certificate and key are properly paired before saving the object.

Supported file formats and import details

You can configure an internal CA object by importing an X.509 v3 CA certificate and private key. You can upload files encoded in one of these supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the private key file is password-protected, you can supply the decryption password. If the certificate and key are encoded in the PEM format, you can also copy and paste the information.

The system only allows you to upload files that contain proper certificate or key information. It also ensures the certificate and key are paired with each other. The system validates the pair before saving the object.



Note If you configure a rule with the **Decrypt - Resign** action, the rule matches traffic based on the referenced internal CA certificate's encryption algorithm type, in addition to any configured rule conditions. You must upload an elliptic curve-based CA certificate to decrypt outgoing traffic encrypted with an elliptic curve-based algorithm, for example.

Import a CA certificate and private key

Use this task when you need to add a new internal CA to your Secure Firewall Management Center for certificate management.

Administrators typically perform this task during initial setup or when updating trusted certificate authorities. It is relevant for those managing PKI infrastructure within Secure Firewall Management Center.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Import CA**.

Step 4 Enter a **Name**.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 Above the **Key** field, click **Browse** to upload a DER or PEM-encoded paired private key file.

Step 7 If the uploaded file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Generate a new CA certificate and private key

Generate a self-signed RSA-based CA certificate and accompanying private key by configuring an internal CA object.

Use this task to establish a trusted CA for your environment. The CA certificate enables secure signing and validation of other certificates and will be valid for ten years, with the **Valid From** date starting one week before creation.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Generate CA**.

Step 4 Enter a **Name**.

Step 5 Enter your identification attributes.

Step 6 Click **Generate self-signed CA**.

Signed certificates

A signed certificate is a digital security credential that

- is obtained from a certificate authority (CA)
- enables configuration of an internal CA object, and
- allows referencing the internal CA object in an SSL rule if it contains a signed certificate.

Certificate configuration steps

You can configure an internal CA object by obtaining a signed certificate from a CA. This process consists of these actions:

- Provide identification information to configure the internal CA object. This action generates an unsigned certificate, produces a paired private key, and creates a certificate signing request (CSR) to a specified CA.
- Upload the signed certificate issued by the CA to the internal CA object. The internal CA object now contains a signed certificate instead of the unsigned certificate.

Only reference an internal CA object in an SSL rule if it contains a signed certificate.

Create an unsigned CA certificate and CSR

Create an unsigned CA certificate and generate a certificate signing request (CSR) to enable secure certificate management in Secure Firewall Management Center.

Perform this task when you need to generate a new unsigned CA certificate and CSR, typically for setting up internal certificate authorities or renewing certificates for secure communications. Always complete this task before uploading a signed certificate issued by a certificate authority.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Generate CA**.

Step 4 Enter a **Name**.

Step 5 Enter the identification attributes.

Step 6 Click **Generate CSR**.

Step 7 Copy the CSR to submit to a CA.

Step 8 Click **OK**.

What to do next

- Upload a signed certificate issued by a CA as described in Upload a signed certificate issued in response to a CSR, on page 2990

Upload a signed certificate issued in response to a CSR

Perform this task after receiving a signed certificate from a certificate authority (CA) in response to a CSR generated by the system. Once uploaded, the certificate is available for referencing in SSL rules.

Before you begin

Obtain the signed certificate file (DER or PEM-encoded X.509 v3 CA certificate) that was issued in response to the CSR.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Edit** ( ) next to the CA object containing the unsigned certificate awaiting the CSR.

Step 4 Click **Install Certificate**.

Step 5 Click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 If the uploaded file is password protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 7 Click **Save** to upload a signed certificate to the CA object.

What to do next

- If an active policy references your object, deploy configuration changes.

CA certificates and private key downloads

A CA certificate and private key download is a security operation that

- enables you to back up or transfer a CA certificate and its paired private key by downloading a file from an internal CA object,
- requires the system to decrypt the private key before creating the downloadable file, and
- uses a password provided by the user to encrypt the downloaded file for secure storage.

Encryption and storage of downloaded CA certificates and private keys

When you download a CA certificate and its private key, the system decrypts the private key before generating the file. Provide a password to encrypt the downloaded file. This helps keep it safe during storage and transfer. The system also encrypts the private key in the CA object with a randomly generated key before saving it to disk.

Caution Store downloaded key information in a secure location.

Caution If you download private keys as part of a system backup, they are decrypted and stored in the unencrypted backup file.

Download a CA certificate and private key

Obtain a CA certificate and private key required for certificate backup, migration, or integration within your secure firewall environment.

You can download CA certificates for both the current domain and ancestor domains. Use this task when you need to obtain the CA certificate and private key for backup, migration, or integration purposes.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Select the internal CA object whose certificate and private key you want to download. Click **Edit** ( ).

Step 4 Click **Download**.

Step 5 Enter an encryption password in the **Password** and **Confirm Password** fields.

Step 6 Click **OK**.

Trusted certificate authority objects

A trusted certificate authority object is a security configuration item that

- represents a CA public key certificate belonging to a trusted CA,
- consists of the object name and CA public key certificate, and
- can be used in SSL policy, realm configurations, and integration with ISE or ISE-PIC.

Usage scenarios for trusted certificate authority objects

You can use trusted certificate authority objects and groups in the following configurations:

- SSL policy: Control traffic encrypted with a certificate signed either by the trusted CA, or any CA within the chain of trust.
- Realm configurations: Establish secure connections to LDAP or Active Directory (AD) servers.
- ISE or ISE-PIC connection: Select trusted certificate authority objects for the pxGrid Server CA and MNT Server CA fields.

After you create the trusted CA object, you can modify the name and add certificate revocation lists (CRL), but cannot modify other object properties. There is no limit on the number of CRLs you can add to an object. If you want to modify a CRL you have uploaded to an object, you must delete the object and recreate it.

Note Adding a CRL to an object has no effect when the object is used in your ISE or ISE-PIC integration configuration.

You cannot delete a trusted CA object that is in use. Additionally, after you edit a trusted CA object that is in use, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

Trusted CA objects

A trusted CA object is a security object that:

- enables configuration of an external CA object by uploading an X.509 v3 CA certificate
- supports uploading files encoded in DER or PEM formats, and
- requires the system to validate the certificate before saving the object.

You can upload a file encoded in one of these supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the file is password-protected, you must supply the decryption password. If the certificate is encoded in the PEM format, you can also copy and paste the certificate information.

You can upload a CA certificate only if the file contains proper certificate information; the system validates the certificate before saving the object.

Add a trusted CA object

Add a trusted CA object so the system can validate certificates issued by a Certificate Authority (CA).

Use this task when you need to import a new trusted CA certificate into the system, such as when enabling secure communications or integrating with external certificate authorities.

Before you begin

- Have the DER or PEM-encoded X.509 v3 CA certificate file available.
- If the file is password-protected, ensure you have the password ready.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Trusted CAs**.

Step 3 Click **Add Trusted CAs**.

Step 4 Enter a **Name**.

Step 5 Click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 If the file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 7 Click **Save**.

The trusted CA object is added and available for use in certificate validation processes.

What to do next

- If an active policy references your object, deploy configuration changes.

Certificate revocation lists in trusted CA objects

A certificate revocation list is a security file that

- can be uploaded to a trusted CA object
- enables control of encrypted traffic based on whether the issuing CA has revoked a certificate, and
- supports files encoded in Distinguished Encoding Rules (DER) or Privacy-enhanced Electronic Mail (PEM) formats.

Certificate revocation list usage in trusted CA objects

You can upload CRLs to a trusted CA object. If you reference that trusted CA object in an SSL policy, you can control encrypted traffic based on whether the CA that issued the session encryption certificate subsequently revoked the certificate.

Supported file formats for CRLs include:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

After you add the CRL, you can view the list of revoked certificates. If you want to modify a CRL you have uploaded to an object, you must delete the object and recreate it.

You can upload only files that contain a proper CRL. There is no limit to the number of CRLs you can add to a trusted CA object. However, you must save the object each time you upload a CRL, before adding another CRL.

Note Adding a CRL to an object has no effect when the object is used in your ISE/ISE-PIC integration configuration.

Add a certificate revocation list to a trusted CA object

Add a certificate revocation list (CRL) to a trusted certificate authority (CA) object. This step ensures that revoked certificates are recognized during certificate validation.

Add a CRL to a trusted CA object to ensure that Secure Firewall Management Center validates client certificates against a list of revoked certificates.

Note Adding a CRL to an object has no effect when the object is used in your ISE/ISE-PIC integration configuration.

Before you begin

Ensure you have a DER or PEM-encoded CRL file available for upload.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Trusted CAs**.

Step 3 Click **Edit** ( ) next to a trusted CA object.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **Add CRL** and upload your DER or PEM-encoded CRL file.

Step 5 Click **OK** to save the changes.

The certificate revocation list is added to the trusted CA object, updating it with the new CRL information.

What to do next

- If an active policy references your object, deploy configuration changes.

External certificate objects

An external certificate object is a certificate management item that

- represents a server public key certificate that does not belong to your organization,
- consists of the object name and certificate, and
- can be used in SSL rules to control traffic that is encrypted with the server certificate. For example, you can upload a self-signed server certificate that you trust, but cannot verify by using a trusted CA certificate.

Supported certificate formats

You can configure an external certificate object by uploading an X.509 v3 server certificate. The supported file formats include:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

The system validates the file before saving the object, ensuring your certificate is correctly processed. If the certificate is encoded in the PEM format, you can also copy and paste the information.

Add external certificate objects

Add external certificate objects to facilitate secure certificate management and integration with external systems.

Use this task when you need to import and manage external X.509 v3 server certificates in Secure Firewall Management Center for PKI integration.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **External Certs**.

Step 3 Click **Add External Cert**.

Step 4 Enter a **Name**.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 server certificate file.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Internal certificate objects

An internal certificate object is a server public key certificate representation that

- belongs to your organization and consists of the object name, public key certificate, and paired private key
- supports SSL rule traffic decryption using the known private key, and
- enables authentication for ISE connections and captive portal configurations.

Internal certificate object usage and configuration

You can use internal certificate objects and groups in:

- your SSL rules to decrypt traffic incoming to one of your organization's servers using the known private key.
- your ISE/ISE-PIC connection. Select an internal certificate object for the **MC Server Certificate** field.
- your captive portal configuration to authenticate the identity of your captive portal device when connecting to users' web browsers. Select an internal certificate object for the **Server Certificate** field.

You can configure an internal certificate object by uploading an X.509 v3 RSA-based or elliptic curve-based server certificate and paired private key. You can upload a file in one of these supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the file is password-protected, you must supply the decryption password. If the certificate and key are encoded in the PEM format, you can also copy and paste the information.

You can upload only files that contain proper certificate or key information, and that are paired with each other. The system validates the pair before saving the object.

After you create the internal certificate object, you can modify the name, but cannot modify other object properties.

You cannot delete an internal certificate object that is in use. Additionally, after you edit an internal certificate object that is in use, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

Add internal certificate objects

Add and manage internal certificate objects for secure handling of server certificates and associated private keys within the system.

Use this task when you need to upload and manage DER or PEM-encoded X.509 v3 server certificates and their corresponding private keys for internal use. This is typically required for secure communication and authentication within your environment.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal Certs**.

Step 3 Click **Add Internal Cert**.

Step 4 Enter a **Name**.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 server certificate file.

Step 6 Above the **Key** field, or click **Browse** to upload a DER or PEM-encoded paired private key file.

Step 7 If the uploaded private key file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 8 Click **Save**.

Certificate enrollment objects

A certificate enrollment object is a PKI configuration item that

- contains the Certification Authority (CA) server information and enrollment parameters required for creating Certificate Signing Requests (CSRs)
- obtains identity certificates from the specified CA, and
- may include certificate revocation information.

Certificate enrollment object reference information

Trustpoints let you manage and track CAs and certificates. A trustpoint is a representation of a CA or identity pair. A trustpoint includes the identity of the CA, CA-specific configuration parameters, and an association with one, enrolled identity certificate.

This section describes how to use and manage certificate enrollment objects in your PKI infrastructure.

- Certificate enrollment objects are used to enroll managed devices into your PKI infrastructure and create trustpoints (CA objects) on devices that support VPN connections.
- To manage certificate enrollment objects, go to **Objects** > **Object Management** > **PKI** > **Certificate Enrollment**. The following information is shown:
  - Existing certificate enrollment objects are listed in the **Name** column. Use the search field (the magnifying glass) to filter the list.
  - The enrollment type of each object is shown in the **Type** column. These enrollment methods can be used:
    - **Self Signed**—The managed device generates its own self signed root certificate.
    - **EST**—Enrollment over Secure Transport is used by the device to obtain an identity certificate from the CA.
    - **SCEP**—(Default) Simple Certificate Enrollment Protocol (SCEP) is used by the device to obtain an identity certificate from the CA.
    - **Manual**—The process of enrolling is carried out manually by the administrator.
    - **PKCS12 File**—Import a PKCS12 file on a Threat Defense managed device that supports VPN connectivity. A PKCS#12, or PFX or P12 file holds the server certificate, any intermediate certificates, and the private key in one encrypted file. Enter the **Passphrase** value for decryption.
  - The **Override** column indicates whether the object allows overrides (a green check mark) or not (a red X). If a number is displayed, it is the number of overrides in place. Use the Override option to customize the object settings for each device that is part of the VPN configuration. Overriding makes each device's trustpoint details unique. Typically the Common Name or Subject is overridden for each device in the VPN configuration. See Object overrides, on page 2939 for details and procedures on overriding objects of any type.
  - **Edit** a previously created certificate enrollment object by clicking on the edit icon (a pencil). Editing can only be done if the enrollment object is not associated with any managed devices. Refer to the adding instructions for editing a certificate enrollment object. Failed enrollment objects can be edited.
  - **Delete** a previously created certificate enrollment object by clicking on the delete icon (a trash can). You cannot delete a certificate enrollment object if it is associated with any managed device.

1. Define parameters for CA authentication and enrollment in a certificate enrollment object. Specify shared parameters and use the override facility to specify unique object settings for different devices.
2. Associate and install this object on each managed device that requires the identity certificate. On the device, it becomes a trustpoint.

   When a certificate enrollment object is associated with and then installed on a device, the process of certificate enrollment starts immediately. The process is automatic for self-signed, SCEP, EST, and PKCS12 file enrollment types, meaning it does not require any additional administrator action. Manual certificate enrollment requires extra administrator action.
3. Specify the created trustpoint in your VPN configuration.

Press (+)**Add Cert Enrollment** to open the **Add Cert Enrollment** dialog and configure a certificate enrollment object, see Add certificate enrollment objects, on page 2999. Then install the certificate on each managed, headend device.

For more information on PKI, digital certificates, and certificate enrollment see PKI infrastructure and digital certificates, on page 1718.

Add certificate enrollment objects

Add certificate enrollment objects to manage device trustpoints and enable secure certificate enrollment for Firewall Threat Defense devices.

- Support multiple enrollment protocols, including SCEP, EST, ACME, Manual, and PKCS12 file import.

Use these objects to enable secure certificate management and device trustpoint enrollment. You must have Admin or Network Admin privileges to do this task.

Procedure

Step 1 Open the **Add Certificate Enrollment** dialog:

- Directly from Object Management: In the **Objects** > **Object Management** > **PKI** > **Certificate Enrollment** from the navigation pane, and click **Add Certificate Enrollment**.
- While configuring a managed device: In the **Devices** > **Certificates** screen, choose **Add** > **Add New Certificate** and click (+) for the **Certificate Enrollment** field.

Step 2 Enter the **Name**. When enrollment is complete, this name is used as the trustpoint name on the managed devices with which it is associated. Click the **CA Information** tab, and then choose the **Enrollment Type**.

- **Self-Signed Certificate**—The managed device, acting as a CA, generates its own self-signed root certificate. No other information is needed in this pane.

  Note
  When enrolling a self-signed certificate you must specify the Common Name (CN) in the certificate parameters.

- **EST**—Enrollment over Secure Transport protocol. Specify the EST information. See EST enrollment configuration options in Certificate Enrollment Object , on page 3001.
- **SCEP**—(Default) Simple Certificate Enrollment Protocol. Specify the SCEP information. See Certificate Enrollment Object SCEP options , on page 3002.
- **Manual**
  - **CA Only**—Check this check box to create only the CA certificate from the selected CA. An identity certificate will not be created for this certificate.

    If you do not select this check box, a CA certificate is not mandatory. You can generate the CSR without having a CA certificate and obtain the identity certificate.
  - **CA Certificate**—Paste the CA certificate in the PEM format in the box. You can also obtain a CA certificate by copying it from another device.

    You can leave this box empty if you choose to generate a CSR without the CA certificate.

- **PKCS12 File**—Import a PKCS12 file on a Firewall Threat Defense managed device that supports VPN connectivity. A PKCS#12, or PFX, file holds a server certificate, intermediate certificates, and a private key in one encrypted file. Enter the **Passphrase** value for decryption.
- **ACME**—ACME protocol is an open and standardized protocol designed to automate the issuance, renewal, and management of SSL and TLS certificates. Management Center communicates with an ACME-enabled CA server and after domain validation, the ACME server issues an SSL or TLS certificate to the device. Refer to Certificate enrollment object options for ACME certificates, on page 3003.

Step 3 **Skip Check for CA flag in basic constraints of the CA Certificate**—Check this check box if you want to skip checking the basic constraints extension and the CA flag in a trustpoint certificate.

Step 4 **Validation Usage**—Choose from the options to validate the certificate during a VPN connection:

- **IPsec Client**—Validate an IPsec client certificate for a site-to-site VPN connection.
- **SSL Client**—Validate an SSL client certificate during a remote access VPN connection attempt.
- **SSL Server**—Select to validate an SSL server certificate, such as a Cisco Umbrella server certificate.

Step 5 (Optional) Click the **Certificate Parameters** tab and specify the certificate contents. Refer to Certificate Enrollment Object certificate parameters in certificate requests , on page 3008.

This information is placed in the certificate and is readable by any party who receives the certificate from the router.

Step 6 (Optional) Click the **Key** tab and specify the Key information. Refer to Certificate Enrollment Object key options , on page 3009.

Step 7 (Optional) Click the **Revocation** tab and specify the revocation options: Refer to Certificate Enrollment Object revocation options , on page 3011.

Step 8 **Allow Overrides** of this object if desired.

When you allow overrides in the PKCS12 certificate enrollment object, update the**Passphrase** for the certificate on the device where you override it.

See Object overrides, on page 2939 for a full description of object overrides.

Step 9 Click **Save**.

What to do next

Associate and install the enrollment object on a device to create a trustpoint on that device.

Related Topics

Install a certificate using self-signed enrollment, on page 3084
Install a certificate using EST enrollment, on page 3084
Install a certificate using SCEP enrollment, on page 3085
Install a certificate using manual enrollment, on page 3086
Install a certificate using a PKCS12 file, on page 3087

Add a certificate enrollment

A valid IdP certificate ensures that your system can securely authenticate users through an identity provider. This is essential for secure communications and compliance.

Procedure

Step 1 Enter the **Name**.

Step 2 Paste the certificate information in the **IdP Certificate** field in PEM format.

Note
If the certificate is dependent on a root or intermediate certificate, you must install the dependent certificates. See Certificates, on page 3079.

Step 3 Click **Save**.

EST enrollment configuration options in Certificate Enrollment Object

This topic describes the EST enrollment options and the required fields for configuring EST certificate enrollment in the Secure Firewall Management Center.

Cloud-Delivered Firewall Management Center navigation path

Objects > **Object Management** > **PKI** > **Certificate Enrollment**. Click **Add Certificate Enrollment** to open the **Add Certificate Enrollment** dialog, and select the **CA Information** tab.

Fields

Enrollment Type—set to **EST**.

Note

- EST enrollment type does not support EdDSA key.
- EST's ability to auto-enroll a device when its certificate expires is not supported.

Enrollment URL—Enter the URL of the CA server to which devices should attempt to enroll.

Use an HTTPS URL in the form of **https://CA_name:port**, where *CA_name* is the host DNS name or IP address of the CA server. The *port* number is mandatory.

Username—The username required to access the CA server.

Password / Confirm Password—The password required to access the CA server.

Fingerprint—Optionally, enter the fingerprint of the CA server certificate. Verifying the fingerprint helps ensure you are connecting to the legitimate CA server and not a substitute. Enter the fingerprint in hexadecimal format. If the fingerprint does not match, the certificate is rejected. Obtain the fingerprint directly from the CA server.

Source Interface—The interface that interacts with the CA server. By default, the diagnostic interface is displayed. To configure a data interface as the source interface, choose the respective security zone or interface group object.

Ignore EST Server Certificate Validations—By default, EST server certificate validation is enabled. Check the check box to ignore Firewall Threat Defense validating EST server certificate.

Certificate Enrollment Object SCEP options

This topic describes the SCEP enrollment options and fields available when adding a certificate enrollment object in Secure Firewall Management Center.

Cloud-Delivered Firewall Management Center navigation path

Objects > **Object Management** > **PKI** > **Certificate Enrollment**. Click **Add Certificate Enrollment** to open the **Add Certificate Enrollment** dialog, and select the **CA Information** tab.

Fields

Enrollment Type—Set to **SCEP**.

Enrollment URL—Enter the URL of the CA server to which devices should attempt to enroll.

Use an HTTP URL in the format **`http://CA_name:port,`** . The CA_name is the DNS name or IP address of the CA server, and the port number is mandatory.



Note If the SCEP Server is referred with hostname/FQDN, configure the DNS Server using FlexConfig object.

If the CA cgi-bin script location at the CA is not the default (/cgi-bin/pkiclient.exe) include the nonstandard script location in the URL, in the form of http://CA_name:port/script_location, where script_location is the full path to the CA scripts.

Challenge Password / Confirm Password—The password used by the CA server to validate the identity of the device. You can obtain the password by contacting the CA server directly or by entering the following address in a web browser: **`http://URLHostName/certsrv/mscep/mscep.dll`**. Because the password is valid for only 60 minutes, deploy it as soon as possible after creation.

Retry Period—The interval (in minutes) between certificate request attempts. Values range from 1 to 60 minutes. The default is 1 minute.

Retry Count—The number of retries that should be made if no certificate is issued on the first request. Value can be 1 to 100. The default is 10.

CA Certificate Source—Specifies how the CA certificate will be obtained.

- **Retrieve Using SCEP** (default and only supported option)—Retrieve the certificate from the CA server using the Simple Certificate Enrollment Process (SCEP). Using SCEP requires a connection between your device and the CA server. Ensure a network route exists between your device and the CA server before beginning enrollment.

Fingerprint—When retrieving the CA certificate using SCEP, you may enter the fingerprint for the CA server. Using the fingerprint to verify the authenticity of the CA server's certificate helps prevent an unauthorized party from substituting a fake certificate in place of the real one. Enter the **Fingerprint** for the CA server in hexadecimal format. If the value you enter does not match the fingerprint on the certificate, the certificate is rejected. Obtain the CA's fingerprint by contacting the server directly, or by opening: **`http://<URLHostName>/certsrv/mscep/mscep.dll`**.

Certificate enrollment object options for ACME certificates

This topic details the ACME options available when configuring certificate enrollment objects, including navigation steps, workflow, and field definitions for ACME certificate enrollment.

Cloud-Delivered Firewall Management Center navigation path

To access ACME certificate enrollment options:

1. Go to **Objects** > **Object Management**.
2. From the navigation pane, choose **PKI** > **Certificate Enrollment**.
3. Click **Add Certificate Enrollment** to open the Add Certificate Enrollment dialog box.
4. Select the **CA Information** tab.

How ACME certificates work

ACME certificate enrollment involves the following steps:

1. A Firewall Threat Defense device requests an ACME certificate for a specific domain or a list of domains through the source interface.
2. The ACME server validates the domain ownership through port 80 of the Firewall Threat Defense device's authentication interface. For domain validation, the Firewall Threat Defense device uses an HTTP-based challenge mechanism (HTTP-01).

Note Firewall Threat Defense briefly opens port 80 during the enrollment challenge-response process to provide only the ACME challenge data; the port closes as soon as the enrollment succeeds or fails.

3. After domain validation, the ACME server issues an SSL or TLS certificate to the device.

Note The stages are repeated for each FQDN in the certificate enrollment request.

ACME Fields and Options

Enrollment Type—Choose **ACME**.

Enrollment URL—Enter the URL of the ACME server.

The default URL is that of the Let's Encrypt service: https://acme-v02.api.letsencrypt.org/directory.

In the **Authentication Protocol** field, **HTTP-01** is the predefined protocol used to validate domain ownership.

Authentication Interface—Choose a security zone or an interface group that has the interface through which the ACME server communicates with the device to verify ownership of the domain. Click + to add a security zone or an interface group.

Source Interface—Choose a security zone or an interface group that has the interface through which the device interacts with the ACME server to request and receive the enrolled ACME certificate. Click + to add a security zone or an interface group.

ACME CA Certificate—Choose a manually installed CA-only certificate that authenticates the ACME server.

Enable Auto Enrollment—Check this check box to enable automatic enrollment of the ACME certificates based on the configured lifetime.

Certificate Lifetime—Enter the percentage of the ACME certificate lifetime after which certificate re-enrollment automatically starts (default value is 70). For example, if the certificate's lifetime is 100 days and this field is set to 80, auto-renewal is triggered on day 81.

Regenerate Key—Check this check box to generate a new key for each ACME enrollment. If you uncheck this check box, the previous key is used for enrollment.

Prerequisites for using ACME certificates

Review these prerequisites for configuring a remote access VPN policy.

General prerequisites

- The Firewall Threat Defense device is Version 10.0 or later.
- Configure DNS in the Firewall Threat Defense platform settings to resolve the domain name of the ACME server.
- Ensure your domain maps to a public IP address. Configure the device interface with this IP address, and set it as the authentication interface in the ACME certificate enrollment.
- Configure the same NTP server for the ACME server and the Firewall Threat Defense device.
- Ensure that you enroll an ACME CA certificate (manually installed CA-only certificate) on the device to authenticate the ACME server.



Note The ACME enrollment will fail if you do not enroll this ACME CA certificate first and add it to the device before configuring any ACME certificate.

For example, if you use Let's Encrypt as the ACME server, you must:

- Get the Internet Security Research Group (ISRG) root certificate from https://letsencrypt.org/certificates/. You can use the root certificate from https://letsencrypt.org/certs/isrgrootx1.pem.txt.

  Sample manual CA certificate for the ACME server from https://letsencrypt.org/certs/isrgrootx1.pem.txt:

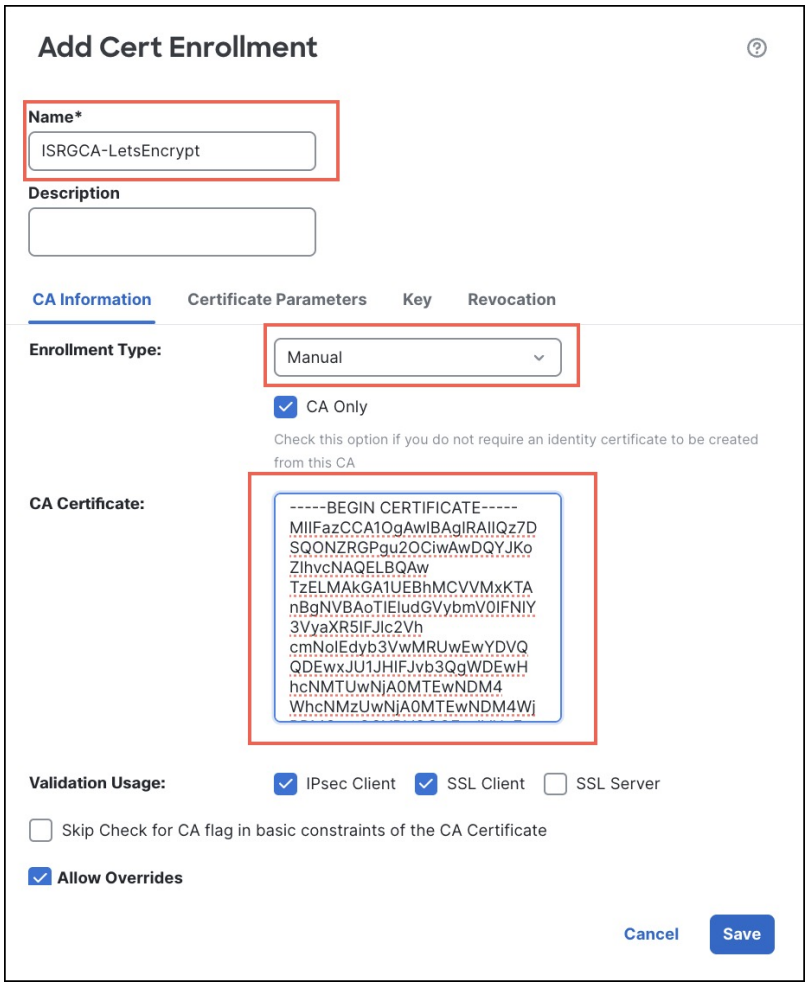
```
-----BEGIN CERTIFICATE-----
MIIFazCCA1OgAwIBAgIRAIIQz7DSQONZRGPgu2OCiwAwDQYJKoZIhvcNAQELBQAw
TzELMAkGA1UEBhMCVVMxKTAnBgNVBAoTIEludGVybmV0IFNlY3VyaXR5IFJlc2Vh
cmNoIEdyb3VwMRUwEwYDVQQDEwxJU1JHIFJvb3QgWDEwHhcNMTUwNjA0MTEwNDM4
WhcNMzUwNjA0MTEwNDM4WjBPMQswCQYDVQQGEwJVUzEpMCcGA1UEChMgSW50ZXJu
ZXQgU2VjdXJpdHkgUmVzZWFyY2ggR3JvdXAxFTATBgNVBAMTDElTUkcgUm9vdCBY
MTCCAiIwDQYJKoZIhvcNAQEBBQADggIPADCCAgoCggIBAK3oJHP0FDfzm54rVygc
h77ct984kIxuPOZXoHj3dcKi/vVqbvYATyjb3miGbESTtrFj/RQSa78f0uoxmyF+
0TM8ukj13Xnfs7j/EvEhmkvBioZxaUpmZmyPfjxwv60pIgbz5MDmgK7iS4+3mX6U
A5/TR5d8mUgjU+g4rk8Kb4Mu0UlXjIB0ttov0DiNewNwIRt18jA8+o+u3dpjq+sW
```

```
T8KOEUt+zwvo/7V3LvSye0rgTBIlDHCNAymg4VMk7BPZ7hm/ELNKjD+Jo2FR3qyH
B5T0Y3HsLuJvW5iB4YlcNHlsdu87kGJ55tukmi8mxdAQ4Q7e2RCOFvu396j3x+UC
B5iPNgiV5+I3lg02dZ77DnKxHZu8A/lJBdiB3QW0KtZB6awBdpUKD9jf1b0SHzUv
KBds0pjBqAlkd25HN7rOrFleaJ1/ctaJxQZBKT5ZPt0m9STJEadao0xAH0ahmbWn
OlFuhjuefXKnEgV4We0+UXgVCwOPjdAvBbI+e0ocS3MFEvzG6uBQE3xDk3SzynTn
jh8BCNAw1FtxNrQHusEwMFxIt4I7mKZ9YIqioymCzLq9gwQbooMDQaHWBfEbwrbw
qHyGO0aoSCqI3Haadr8faqU9GY/rOPNk3sgrDQoo//fb4hVC1CLQJ13hef4Y53CI
rU7m2Ys6xt0nUW7/vGT1M0NPAgMBAAGjQjBAMA4GA1UdDwEB/wQEAwIBBjAPBgNV
HRMBAf8EBTADAQH/MB0GA1UdDgQWBBR5tFnme7bl5AFzgAiIyBpY9umbbjANBgkq
hkiG9w0BAQsFAAOCAgEAVR9YqbyyqFDQDLHYGmkgJykIrGF1XIpu+ILlaS/V9lZL
ubhzEFnTIZd+50xx+7LSYK05qAvqFyFWhfFQDlnrzuBZ6brJFe+GnY+EgPbk6ZGQ
3BebYhtF8GaV0nxvwuo77x/Py9auJ/GpsMiu/X1+mvoiBOv/2X/qkSsisRcOj/KK
NFtY2PwByVS5uCbMiogziUwthDyC3+6WVwW6LLv3xLfHTjuCvjHIInNzktHCgKQ5
ORAzI4JMPJ+GslWYHb4phowim57iaztXOoJwTdwJx4nLCgdNbOhdjsnvzqvHu7Ur
TkXWStAmzOVyyghqpZXjFaH3pO3JLF+l+/+sKAIuvtd7u+Nxe5AW0wdeRlN8NwdC
jNPElpzVmbUq4JUagEiuTDkHzsxHpFKVK7q4+63SM1N95R1NbdWhscdCb+ZAJzVc
oyi3B43njTOQ5yOf+1CceWxG1bQVs5ZufpsMljq4Ui0/1lvh+wjChP4kqKOJ2qxq
4RgqsahDYVvTH9w7jXbyLeiNdd8XM2w9U/t7y0Ff/9yi0GE44Za4rF2LN9d11TPA
mRGunUHBcnWEvgJBQl9nJEiU0Zsnvgc/ubhPgXRR4Xq37Z0j4r7g1SgEEzwxA57d
emyPxgcYxn/eR44/KJ4EBs+lVDR3veyJm+kXQ99b21/+jh5Xos1AnX5iItreGCc=
-----END CERTIFICATE-----
```

- Enroll it as a manual CA-only certificate on the device. If you configure object overrides for any device, ensure that you enroll an ACME CA certificate on that device too.

  To enroll this ACME CA certificate on the device:

  1. Choose **Objects** > **Object Management** > **PKI** > **Certificate Enrollment**.
  2. Click **Add Certificate Enrollment**.
  3. Configure the name and description for the certificate.
  4. In the **CA Information** tab, configure these parameters:
     a. From the **Enrollment Type** drop-down list, choose **Manual**.
     b. Check the **CA Only** check box.
     c. In the CA Certificate text box, paste the manual CA certificate for the ACME server. In our example, we use the root certificate from Let's Encrypt (https://letsencrypt.org/certs/isrgrootx1.pem.txt).

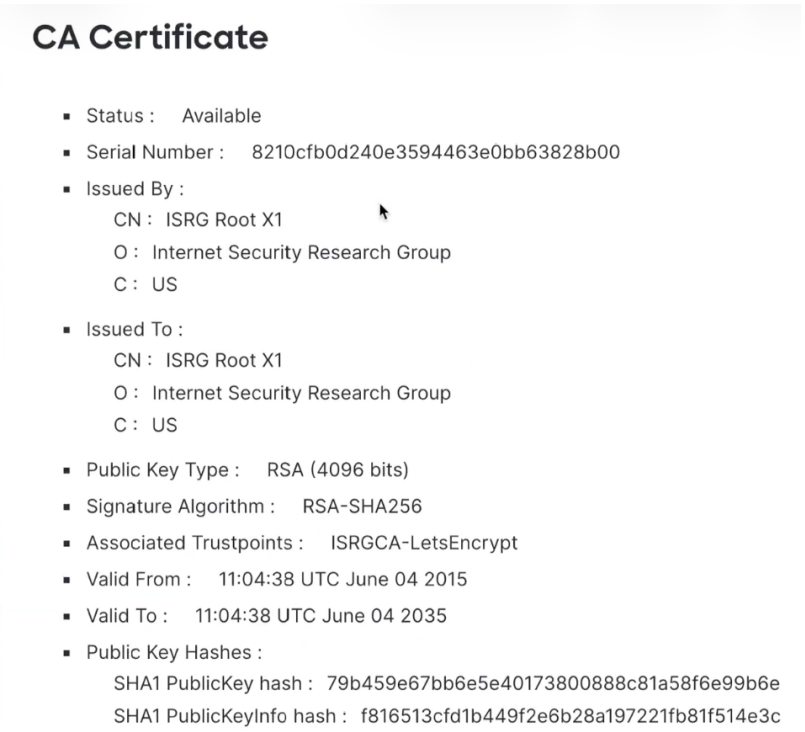


d. Click **Save**.

Note This ACME CA certificate is listed in the **Certificate Enrollment** page and not under **Trusted CAs**.

- Attach this ACME CA certificate to the device:
  1. Choose **Devices** > **Certificates**, and click **Add**.
  2. From the **Device** drop-down list, choose a Firewall Threat Defense device.
  3. From the **Cert Enrollment** drop-down list, choose the ACME CA certificate.
  4. Click **Add**.

Figure 398: Details of the manual CA certificate



Prerequisites for ACME server

- Access to an ACME server such as Let's Encrypt, or any other public or on-prem ACME server.
- The ACME server must be reachable from the Firewall Threat Defense device.
- The ACME server must be able to validate the domain name and the alternate FQDNs.
- If the authentication and source interfaces of the Firewall Threat Defense device differ, ensure that the ACME server is reachable from the source interface of the Firewall Threat Defense device.
    - Authentication interface of the Firewall Threat Defense device is the interface through which the ACME server communicates with the device to verify ownership of the domain.

      Source interface of the Firewall Threat Defense device is the interface through which the device interacts with the ACME server to request and receive the enrolled ACME certificate.
- Firewall Threat Defense device does not support ACME wildcard certificates.

Prerequisite for VPN load balancing

- When configuring an ACME enrollment object for a VPN load balancing group, include both director and member FQDNs in the Alternate FQDN field.

ACME certificate limitations and compatibility considerations

ACME certificates come with several limitations and compatibility requirements. Review these restrictions and supported features before deploying ACME certificates in your environment.

- ACME certificates do not support:
  - Site-to-site VPN
  - Management interface in the converged mode
  - DNS authentication (DNS-01); only HTTP-01 is supported
  - Domain override
  - Wildcard certificates (Certificates that secure a domain and multiple subdomains using a wildcard character (*) in the domain name field)
  - Clustering
- Supported key sizes:
  - RSA keys: 2048, 3072, and 4096 bits
  - ECDSA keys: 256, 384, and 521 bits
- ACME enrolment is not compatible with control plane ACLs.

  When using Let's Encrypt with control plane ACLs:

  1. Disable the ACL before the ACME enrollment to allow port 80 access.
  2. Enroll the ACME certificate.
  3. Verify the enrollment.
  4. Re-enable the ACL.

  When using ACME servers other than Let's Encrypt with control plane ACLs, include the server's FQDN in the ACL.

Certificate Enrollment Object certificate parameters in certificate requests

When sending a certificate request to a CA server, you can specify additional information that is placed into the certificate and is viewable by recipients. All information should follow the standard LDAP X.500 format.

Cloud-Delivered Firewall Management Center navigation path

Objects > **Object Management** > **PKI** > **Certificate Enrollment**. Click **Add Certificate Enrollment** to open the **Add Certificate Enrollment** dialog box, and select the **Certificate Parameters** tab.

Fields

Field	Description
Include FQDN	Whether to include the device's fully qualified domain name (FQDN) in the certificate request. You can select from these options: • **Use Device Hostname as FQDN** • **Don't use FQDN in certificate** • **Custom FQDN**—Select this and then specify it in the **Custom FQDN** field that displays.
Include Device's IP Address	Specify the interface whose IP address you want to include in the certificate request.
Common Name (CN)	Specify the X.500 common name to include in the certificate. **Note** When enrolling a self-signed certificate you must specify the Common Name (CN) in the certificate parameters.
Organization Unit (OU)	Specify the name of the organization unit (for example, a department name) to include in the certificate.
Organization (O)	Specify the organization or company name to include in the certificate.
Locality (L)	Specify the state or province to include in the certificate.
State (ST)	The state or province to include in the certificate.
County Code (C)	The country to include in the certificate. These codes conform to ISO 3166 country abbreviations, for example "US" for the United States of America.
Email (E)	Specify the email address to include in the certificate.
Include Device's Serial Number	Specify whether to include the device's serial number in the certificate. The CA uses the serial number to authenticate certificates or to later associate a certificate with a particular device. If you are unsure, include the serial number, because it is useful for debugging purposes.

Certificate Enrollment Object key options

This topic describes the available key options and their settings when adding certificate enrollment and selecting the Key tab in the Certificate Enrollment Object dialog box.

Cloud-Delivered Firewall Management Center navigation path

Objects > **Object Management** > **PKI** > **Certificate Enrollment**, and click **Add Certificate Enrollment** to open the **Add Certificate Enrollment** dialog box, and select the **Key** tab.

Fields

- **Key Type**: Choose RSA, ECDSA, EdDSA.

Note

- EdDSA is not supported for EST enrollment type.
- EdDSA is supported only in Site-to-Site VPN topologies.
- EdDSA is not supported as an identity certificate for Remote Access VPN.

- **Key Name**: Specifies the name of the key pair. If the key pair already exists, this field refers to the existing name. If not, it defines the name for a new key pair to be generated during enrollment. If no name is specified, the fully qualified domain name (FQDN) key pair is used instead.
- **Key Size**: When generating a new key pair, defines the desired key size (modulus) in bits. The recommended size is 2048 bits. Larger modulus sizes offer more security but require more time (a minute or more when larger than 512 bits) to generate and process.

Important

- On Cloud-Delivered Firewall Management Center and Firewall Threat Defense Versions 7.0 and higher, you cannot enroll certificates with RSA key sizes smaller than 2048 bits and keys using SHA-1 with the RSA Encryption algorithm.
- You can use the PKI Enrollment of Certificates with Weak-Crypto to allow certificates with SHA-1 and smaller key size.
- You cannot generate RSA keys with sizes smaller than 2048 bits for Firewall Threat Defense 7.0, even when you enable the weak-crypto option.

- **Advanced Settings**: Select **Ignore IPsec Key Usage** if you do not want to validate values in the key usage and extended key usage extensions of IPsec remote client certificates. You can suppress key usage checking on IPsec client certificates. By default, this option is not enabled.

Note For site-to-site VPN connection, if you use a Windows Certificate Authority (CA), the default Application Policies extension is **IP security IKE intermediate**. If you are using this default setting, you must select the **Ignore IPsec Key Usage** option for the object you select. Otherwise, the endpoints cannot complete the site-to-site VPN connection.

Enroll PKI certificates with weak-crypto

Certificates that use weak cryptography (such as SHA-1 signatures or RSA keys smaller than 2048 bits) have restricted support on Cloud-Delivered Firewall Management Center and Firewall Threat Defense devices. This reference details the conditions under which weak-crypto certificates may be enrolled and the required steps for enabling their use.

SHA-1 hashing signature algorithm, and RSA key sizes that are smaller than 2048 bits for certification are not supported on Cloud-Delivered Firewall Management Center and Firewall Threat Defense Version 7.0 and later. You cannot enroll certificates with RSA key sizes that are smaller than 2048 bits.

To override these restrictions on Cloud-Delivered Firewall Management Center 7.0 managing Firewall Threat Defenses running Versions earlier than 7.0, use the enable weak-crypto option on the Firewall Threat Defense. Do not permit weak-crypto keys, because, such keys are not as secure as the ones with larger key sizes.

Note Firewall Threat Defense 7.0 or later does not support generating RSA keys with sizes smaller than 2048 bits even when you permit weak-crypto.

To enable weak-crypto on the device, navigate to the **Devices** > **Certificates** page. Click the **Enable Weak-Crypto**() button provided against the Firewall Threat Defense device. When the weak-crypto option is enabled, the button changes to . By default, the weak-crypto option is disabled.

Note When a certificate enrollment fails due to weak cipher usage, the Cloud-Delivered Firewall Management Center displays a warning message prompting you to enable the weak-crypto option. Similarly, when you turn on the enable weak-crypto button, the Cloud-Delivered Firewall Management Center displays a warning message before enabling weak-crypto configuration on the device.

Upgrade earlier versions to Firewall Threat Defense 7.0

When you upgrade to Firewall Threat Defense 7.0, the existing certificate configurations are retained. However, if those certificates have RSA keys smaller than 2048 bits and use SHA-1 encryption algorithm, they cannot be used to establish VPN connections. You must either procure a certificate with RSA key size of at least 2048 bits or enable the permit weak-crypto option for VPN connections.

Certificate Enrollment Object revocation options

Specify whether to check the revocation status of a certificate by choosing and configuring the method. Revocation checking is off by default. The system does not check either method (CRL or OCSP).

Cloud-Delivered Firewall Management Center navigation path

Objects > **Object Management** > **PKI** > **Certificate Enrollment**, and click () **Add Cert Enrollment** to open the **Add Cert Enrollment** dialog, and select the **Revocation** tab.

Fields

- **Enable Certificate Revocation Lists**—Check to enable CRL checking.
  - **Use CRL distribution point from the certificate**—Check to obtain the revocation lists distribution URL from the certificate.
  - **Use static URL configured**—Check this to add a static, predefined distribution URL for revocation lists. Then add the URLs.

    CRL Server URLs—Specify the URL of the LDAP server from which the CRL can be downloaded.

    URLs must start with **`http://`**. Include a port number in the URL. Enclose IPv6 addresses in square brackets, for example: *http://[0:0:0:0:0.18:0a01:7c16]*.

- **Enable Online Certificate Status Protocol (OCSP)**—Check to enable OCSP checking.

OCSP Server URL—The URL of the OCSP server checking for revocation if you require OCSP checks.

URLs must start with **`http://`**. Enclose IPv6 addresses in square brackets, for example: *http://**[**0:0:0:0:0.18:0a01:7c16**]***.

- **Consider the certificate valid if revocation information cannot be reached**—Checked by default. Uncheck if you do not want to allow this.

Configure a policy list

Configure a policy list to group and manage multiple match statements for use in route maps. This allows for efficient policy management and flexible routing decisions.

Use the Configure Policy List page to create, copy, and edit policy list objects. You can create policy list objects to use when you are configuring route maps. When a policy list is referenced within a route map, all of the match statements in that policy list are evaluated and processed. With a route map, you can configure two or more policy lists. A policy list can coexist with other match and set statements that are configured in the same route map but outside the policy list. When multiple policy lists perform matching within a route map entry, only the incoming attribute is matched by all policy lists.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** > **Policy List**.

Step 2 Click **Add Policy List**. Specify these values:

a. Enter a name for the policy list object in the **Name** field. Object names are not case-sensitive.

b. Select whether to allow or block access for matching conditions from the **Action** drop-down list.

Step 3 Enter the details in each of these tabs:

Tabs	Details
Interface	Use this tab to distribute routes that have their next hop out of one of the interfaces specified. In the **Zones/Interfaces** list, add the zones that contain the interfaces through which the device communicates with the management station. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zone/Interface** list and click **Add**. The host will be configured on a device only if the device includes the selected interfaces or zones.
Address	Use this tab to redistribute any routes that have a destination address that is permitted by a standard access list or prefix list. Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.

Tabs	Details
Next Hop	Use this tab to redistribute any routes that have a next hop router address passed by one of the access lists or prefix lists specified. Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.
Route Source	Use this tab to redistribute routes that have been advertised by routers and access servers at the address specified by the access lists or prefix list. Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.
AS Path	Use this tab to match a BGP autonomous system path. If you specify more than one AS path, then the route can match either AS path.
Community Rule	Use this tab to enable matching of the BGP community or extended community with the specified community list objects or the extended community list objects respectively. If you specify more than one rule, the routes are verified against the rules until a matching permit or deny is met. **a.** To specify a community list to the rule, click **Edit** ( ) given in the **Selected Community List** field. The community lists appear under **Available Community List**. Select the required list, click **Add**, and then click **Ok**. To enable matching the BGP community exactly with the specified community, check the **Match the specified community exactly** check box. **b.** To add the extended community list, click **Edit** ( ) given in the **Selected Extended Community List** field. The extended community lists appear under the **Available Extended Community List**. Select the required list, click **Add**, and then click **Ok**. **Note** The extended community lists are applicable only for configuring import or export of routes.
Metric & tag	Use this tab to match the metric and security group tag of a route. **a.** Enter the metric values to use for matching in the **Metric** field. You can enter multiple values, separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295. **b.** Enter the tag values to use for matching in the **Tag** field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.

Step 4 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 5 Click **Save**.

Ports

A port object is a network configuration element that

- represents different protocols, such as TCP, UDP, ICMP, and others, in distinct ways
- can specify protocol numbers, port numbers or ranges, and optional type and code for certain protocols, and
- is used to define access and monitoring rules within the system's web interface.

Port object representations and usage

Port objects represent protocols differently depending on their type.

- **TCP and UDP**: Port objects represent these transport layer protocols by specifying the protocol number, and may include an associated port or port range. For example: `TCP(6)/22`.
- **ICMP and ICMPv6 (IPv6-ICMP)**: Port objects represent these Internet layer protocols by specifying the protocol number and optional type and code. For example: `ICMP(1):3:3`.

  Restrict an ICMP or IPV6-ICMP port object by type and, if needed, code. Learn more about ICMP types and codes at:

  - ICMP Types and Codes: http://www.iana.org/assignments/ICMP-parameters/ICMP-parameters.xml
  - ICMPv6 Types and Codes: http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xml
- **Other protocols**: Port objects can represent protocols that do not use ports.

The system provides default port objects for well-known ports. Default objects are read-only and always available for use. You can create custom port objects or groups in addition to the default ones.

Port objects and groups can be used throughout the system's web interface in access control policies, identity rules, network discovery rules, port variables, and event searches. For example, you can configure a network discovery policy to exclude monitoring ports used by a custom client if it generates excessive or misleading events.

Guidelines for using port objects

- You cannot add any protocol other than TCP or UDP for source port conditions in access control rules.
- Set source and destination port conditions in a rule using only one transport protocol.
- Add only supported protocols to a port object group used in a source port condition so that the rule takes effect on the managed device when you deploy the configuration.
- To add a destination port, create a port object using only TCP or only UDP ports for the source port condition in a rule. Similarly, to add a source port, create a port object using only TCP or only UDP ports as the destination port condition.

Create port objects

Create port objects to specify protocol and port information for network traffic management in Secure Firewall Management Center.

Use this task to add or modify port objects in access control or other security policies within Secure Firewall Management Center. Port objects allow you to specify which network traffic is allowed or blocked based on protocol and port criteria. They are critical when you set up or update security rules.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Port** from the object types list.

Step 3 Choose **Add Object** from the **Add Port** drop-down list.

Alternatively, clone an existing port object and edit its parameters to create a new port object. Click the **Clone** icon on the existing port object that you want to clone.

Step 4 Enter a **Name**.

Step 5 Choose a **Protocol**.

Step 6 Depending on the protocol you chose, constrain by **Port**, or choose an ICMP **Type** and **Code**.

You can enter ports from `1` to `65535`. Use a hyphen to specify a port range. You must constrain the object by port if you chose to match **All** protocols, using the **Other** drop-down list.

Step 7 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.
- If you want to add override values to this object, expand the Override section and click **Add**; see Add object overrides, on page 2942.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Import port objects

To import port objects, refer to Import objects, on page 2932.

Prefix lists

A prefix list is a route filtering object that

- defines IPv4 and IPv6 address ranges

- is used when configuring route maps, policy maps, OSPF filtering, or BGP neighbor filtering, and
- enables control over which routes are permitted or denied.

Configure an IPv6 prefix list

Use IPv6 prefix lists in Firewall Threat Defense devices to specify which IPv6 routes are allowed or blocked during redistribution. This enhances routing policy control for protocols such as OSPF and BGP.

Procedure

Step 1 Select **Objects** > **Object Management** > **Prefix Lists** > **IPv6 Prefix List** from the table of contents.

Step 2 Click **Add Prefix List**.

a. Enter a name for the prefix list object in the **Name** field

b. Click **Add**.

Step 3 Select the appropriate action, Allow or Block, from the **Action** drop-down list to indicate the redistribution access.

Step 4 Enter a unique number that indicates the position a new prefix list entry will have in the list of prefix list entries already configured for this object, in the **Sequence No.** field. If left blank, the sequence number will default to five more than the largest sequence number currently in use.

Step 5 Specify the IPv6 address in the IP address/mask length format in the **IP address** field. The mask length must be a valid value between 1 and 128.

Step 6 Enter the minimum prefix length in the **Minimum Prefix Length** field and the maximum prefix length in the **Maximum Prefix Length** field.

The minimum prefix length value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.

If the Minimum Prefix Length is present, the maximum prefix length value must be greater than or equal to it. If the Minimum Prefix Length is not specified, the maximum value must be greater than the mask length.

Step 7 Click **Add**.

Step 8 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 9 Click **Save**.

Configure an IPv4 prefix list

Use the Configure IPv4 Prefix list page to manage prefix list objects. These objects can be used when configuring route maps, policy maps, OSPF filtering, or BGP neighbor filtering.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** > **Prefix Lists** > **IPv4 Prefix List** from the table of contents.

Step 2 Click **Add Prefix List**.

a. Enter a name for the prefix list object in the **Name** field

b. Click **Add**.

Step 3 Select the appropriate action, Allow or Block, from the **Action** drop-down list to indicate the redistribution access.

Step 4 Enter a unique number in the **Sequence No.** field to specify the position of the new prefix list entry. If the field is left blank, the system assigns a sequence number that is five higher than the largest one currently in use.

Step 5 Specify the IPv4 address in the IP address/mask length format in the **IP address** field. The mask length must be a valid value between 1 and 32.

Step 6 Enter the minimum prefix length in the **Minimum Prefix Length** field and the maximum prefix length in the **Maximum Prefix Length** field.

The minimum prefix length value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.

If the Minimum Prefix Length is present, the maximum prefix length value must be greater than or equal to it. If the Minimum Prefix Length is not specified, the maximum value must be greater than the mask length.

Step 7 Click **Add**.

Step 8 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 9 Click **Save**.

Configure a route map

Route maps are used to manage route redistribution between routing protocols and to control default route generation. Route maps can specify which routes are permitted or denied and apply rules for matching route attributes.

You can use this object with Firewall Threat Defense devices.

Configure a route map to create or edit route map entries in a Route Map object.

Before you begin

A Route Map can use one or more of these objects. It is not mandatory to add all of them. Create and use any of these objects as required to configure your route map.

- Add ACLs.
- Add Prefix Lists.
- Add AS Path.

- Add Community Lists.
- Add Extended Community Lists.



Note The extended community lists apply only to configuring import or export of routes.

- Add Policy Lists.

Procedure

Step 1 Select **Objects** > **Object Management** > **Route Map**.

Step 2 Click **Add Route Map**.

Step 3 Click **Add** on the **New Route Map Object** window.

Step 4 In the **Sequence No.** field, enter a number, from 0 through 65535, that indicates the position a new route map entry has in the list of route maps entries already configured for this route map object.

Note
We recommend that you number clauses in intervals of at least 10 to reserve numbering space in case you want to insert clauses in the future.

Step 5 Select Allow or Block from the **Redistribution** drop-down list, to indicate the redistribution access.

Step 6 Click the **Match Clauses** tab to match routes or traffic based on the following criteria, which you select in the table of contents:

- **Security Zones** —Match traffic based on the (ingress/egress) interfaces. You can select zones and add them, or type in interface names and add them.
- **IPv4**— Match IPv4 (routes/traffic) based on the following criteria; select the tab to define the criteria.
  - **a.** Click the **Address** tab to match routes based on the route address. For IPv4 addresses, choose whether to use an Access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
  - **b.** Click the **Next Hop** tab to match routes based on the next hop address of a route. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
  - **c.** Click the **Route Source** tab to match routes based on the advertising source address of the route. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
- **IPv6** —Match IPv6 (routes/traffic) based on the route address, next-hop address or advertising source address of route.
- **BGP**—Match BGP (routes/traffic) based on the following criteria; select the tab to define the criteria.
  - **a.** Click the **AS Path** tab to enable matching the BGP autonomous system path access list with the specified path access list. If you specify more than one path access list, then the route can match either path access list.

b. Click the **Community List** tab to enable matching of the BGP community or extended community with the specified community list objects or the extended community list objects respectively.

- To specify a community list to the rule, click **Edit** ( ) given in the **Selected Community List** field. The community lists appears under **Available Community List**. Select the required list, click **Add**, and then click **Ok**. For information on how to create community list objects, see Configure a community list, on page 2959
- To add the extended community list, click **Edit** ( ) given in the **Selected Extended Community List** field. The extended community lists appears under the **Available Extended Community List**. Select the required list, click **Add**, and then click **Ok**. For information on how to create extended community list objects, see Configure an extended community list, on page 2960.

To enable matching the BGP community exactly with the specified community list objects, check the **Match the specified community exactly** check box. This option is not applicable for the extended community list.

Note
If you specify more than one rule, the routes are verified against the rules until a matching permit or deny condition is met. Any route that does not match at least one Match community will not be advertised for outbound route maps.

c. Click the **Policy List** tab to configure a route map to evaluate and process a BGP policy. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only.

- **Others** —Match routes or traffic based on the following criteria.

a. Enter the metric values to use for matching in the **Metric Route Value** field, to enable matching the metric of a route. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.

b. Enter the tag values to use for matching in the **Tag Values** field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.

c. Check the appropriate **Route Type** option to enable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. You can choose more than one route type from the list.

Step 7 Click the **Set Clauses** tab to set routes/traffic based on the following criteria, which you select in the table of contents:

- **Metric Values**—Set either Bandwidth, all of the values or none of the values.

a. Enter a metric value or bandwidth in Kbits per second in the **Bandwidth** field. Valid values are an integer value in the range from 0 to 4294967295.

b. Select to specify the type of metric for the destination routing protocol, from the **Metric Type** drop-down list. Valid values are: internal, type-1, or type-2.

- **BGP Clauses** —Set BGP routes based on the following criteria; select the tab to define the criteria.

a. Click the **AS Path** tab to modify an autonomous system path for BGP routes.

1. Enter an AS path number in the **Prepend AS Path** field to prepend an arbitrary autonomous system path string to BGP routes. Usually the local AS number is prepended multiple times, increasing the autonomous system path length. If you specify more than one AS path number then the route can prepend either AS number.
2. Enter an AS path number in the **Prepend Last AS to AS Path** field to prepend the AS path with the last AS number. Enter a value for the AS number from 1 to 10.
3. Check the **Convert route tag into AS path** check box to convert the tag of a route into an autonomous system path.

b. Click the **Community List** tab to set the community attributes:

Under **Specific Community**:

1. Click the **None** radio button, to remove the community attribute from the prefixes that pass the route map.
2. Click the **Specific Community** radio button, to enter a community number, if applicable. Valid values are from 1 to 4294967295.
3. Check the **Add to existing communities** check box, to add the community to the already existing communities.
4. Select the **Internet**, **No-Advertise**, or **No-Export** check-boxes to use one of the well-known communities.

Under **Specific Extended Community**, in the **Route Target** field, enter the route target number in *ASN:nn* format:

- You can enter values that ranges from 1:1 to 65534:65535.

  You can add a single route target or a set of route targets separated by commas in a single entry. For example, *1:2,1:4,1:6*.

- You can have a maximum of 8 route targets in an entry.
- You cannot have redundant route target entries across route maps.

c. Click the **Others** tab to set additional attributes.

1. Check the **Set Automatic Tag** check-box to automatically compute the tag value.
2. Enter a preference value for the autonomous system path in the **Set Local Preference** field. Enter a value between 0 and 4294967295.
3. Enter a BGP weight for the routing table in the **Set Weight** field. Enter a value between 0 and 65535.
4. Select to specify the BGP origin code. Valid values are **Local IGP** and **Incomplete**.
5. In the IPv4 Settings section, specify a next hop IPv4 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv4 address then the packets can output at either IP address.

   Select to specify an IPv4 prefix list in the **Prefix List** drop-down list.

6. In the IPv6 Settings section, specify a next hop IPv6 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv6 address, then the packets can output at any of the IP addresses.

   Select to specify an IPv6 prefix in the **Prefix List** drop-down list.

Step 8 Click **Add**.

If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.

Step 9 Click **Save**.

After completing these steps, the route map is configured and ready for use in route redistribution and policy enforcement. The route map entries and associated match and set clauses are applied AS specified.

What to do next

Review the route map configuration to ensure that all required match and set clauses are correctly defined. Make any necessary adjustments to optimize routing policies.

- Monitor route redistribution and verify that routes are processed according to the configured route map.

Security intelligence

Security Intelligence is a traffic filtering mechanism that

- uses collections of IP addresses, domain names, and URLs to quickly filter traffic that matches entries
- requires the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types), and
- operates through lists and feeds that are grouped into DNS, Network, and URL categories.

Security intelligence components

Security Intelligence *lists* and *feeds* are collections of IP addresses, domain names. You can use these to quickly filter traffic that matches an entry on a list or feed.

- A list is a static collection that you manage manually.
- A feed is a dynamic collection that updates on an interval over HTTP or HTTPS.

Security Intelligence lists/feeds are grouped into:

- DNS (Domain names)
- Network (IP addresses)
- URLs

System-provided feeds:

Cisco provides these feeds as Security Intelligence objects:

- Security Intelligence feeds updated regularly with the latest threat intelligence from Talos:
  - Cisco-DNS-and-URL-Intelligence-Feed (under DNS Lists and Feeds)
  - Cisco-Intelligence-Feed (for IP addresses, under Network Lists and Feeds)

  You cannot delete the system-provided feeds, but you can change the frequency of (or disable) their updates.
- Cisco-TID-Feed (under Network Lists and Feeds)

  This feed is not used in the Security Intelligence tab of the access control policy.

  Instead, you must enable and configure Secure Firewall Threat Intelligence Director to use this feed. The feed is a collection of TID observables data.

  Use this object to configure the frequency of data publication to TID elements.

Predefined global Block and Do Not Block lists:

The system includes predefined global Block lists and Do Not Block lists for domains (DNS), IP addresses (Networks), and URLs.

These lists are empty until you populate them. To build these lists, see Global and domain security intelligence lists, on page 3023.

By default, access control and DNS policies use these lists as part of Security Intelligence.

Custom feeds:

You can use third-party feeds or use a custom internal feed to easily maintain an enterprise-wide Block list in a large deployment with multiple Cloud-Delivered Firewall Management Center appliances.

See Custom security intelligence feeds, on page 3029.

Custom lists:

Custom lists can augment and fine-tune feeds and the global lists.

See Custom Security Intelligence lists, on page 3031.

Security Intelligence lists and feeds usage locations:

- IP address and address blocks—Use Block and Do Not Block lists in access control policies, as part of Security Intelligence.
- Domain Names—Use Block and Do Not Block lists in DNS policies, as part of Security Intelligence.
- URLs—Use Block and Do Not Block lists in access control and QoS rules, whose analysis and traffic handling phases occur after Security Intelligence.

Modify security intelligence objects

This topic explains how to add or delete entries on a Block list, Do Not Block list, feed, or sinkhole object. It also describes the edit capabilities and redeployment requirements for each object type.

The table summarizes the editing options and redeployment requirements for various security intelligence object types.

Object Type	Edit Capabilities	Requires Redeploy After Edit?
Custom Block and Do Not Block lists	Upload new and replacement lists using the object manager.	No
Default (but custom-populated) Block lists and Do Not Block lists: Global, descendant, and domain-specific	Add entries using the context menu or delete entries using the object manager.	No
System-provided Intelligence Feeds	Disable or change update frequency using the object manager.	No
Custom feeds	Fully modify using the object manager.	No
Sinkhole	Fully modify using the object manager.	Yes

Global and domain security intelligence lists

A security intelligence list helps you manage firewall policies by

- lets you manage lists of connections to block or exempt,
- applies your policies consistently across global and domain contexts, and
- enables further threat evaluation without the need to redeploy firewall policies.

Security intelligence list reference information

The Cloud-Delivered Firewall Management Center includes Global Block and Do Not Block lists, so you can use Security Intelligence to consistently block or exempt connections, allowing evaluation by other threat detection processes you configure.

Access control and DNS policies use these Global lists by default, and they apply to all security zones. If you want, you can turn these lists off for individual policies.

- Global Block list: Lets you block specified connections in all security zones.
- Do-Not-Block list: Lets you exempt specified connections from blocking so you can further evaluate threats.



Note These options apply only to Security Intelligence. Security Intelligence does not block traffic already fastpathed. Adding an item to a Security Intelligence Do Not Block list does not mean matching traffic is trusted or fastpathed. For more information, see Security intelligence, on page 2027.

Global and domain security intelligence lists example

For example, if you notice a set of routable IP addresses in intrusion events associated with exploit attempts, you can immediately block those IP addresses. Although it may take a few minutes for your changes to propagate, you do not have to redeploy.

Security intelligence lists and multitenancy

A security intelligence list is a multitenancy feature that

- enables the creation of block or do not block lists that apply to specific subdomains,
- supports descendant domain lists that aggregate the lists of a domain's descendants, and
- allows administrators at or above a domain to populate lists, while only domain-specific administrators can remove items.

Domain lists and descendant domain lists

Domain lists are block or do not block lists whose contents apply to a particular subdomain. The global lists are domain lists for the global domain. Each subdomain has its own named lists, and their contents apply only to that subdomain.

For example, a subdomain named Company A owns these lists:

- Domain Block list - Company A and Domain Do Not Block list - Company A
- Domain Block list for DNS - Company A, Domain Do Not Block list for DNS - Company A
- Domain Block list for URL - Company A, Domain Do Not Block list for URL - Company A

Any administrator at or above the current domain can populate these lists. You can use the context menu to add an item to the block or do not block list in the current and all descendant domains. However, only an administrator in the associated domain can remove an item from a domain list.

For example, a global administrator could add the same IP address to the block list in the global domain and Company A's domain, but not in Company B's domain. This action would add the same IP address to:

- Global Block list (removable only by global administrators)
- Domain Block list - Company A (removable only by Company A administrators)

A descendant domain list is a do not block list or block list that aggregates the domain lists of the current domain's descendants. Leaf domains do not have descendant domain lists.

Descendant domain lists are useful because a higher-level domain administrator can enforce general security intelligence settings, while subdomain users can add items to a block or do not block list in their own deployment.

For example, the global domain has these descendant domain lists:

- Descendant Block lists - Global, Descendant Do Not Block lists - Global
- Descendant Block lists for DNS - Global, Descendant Do Not Block lists for DNS - Global
- Descendant Block lists for URL - Global, Descendant Do Not Block lists for URL - Global

Note Descendant domain lists do not appear in the object manager because they are symbolic aggregations, not hand-populated lists. They appear where you can use them: in access control and DNS policies.

Delete entries from global security intelligence lists

Update your security intelligence lists by removing entries that should no longer be blocked or allowed. This helps you maintain accurate and effective security intelligence lists.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Security Intelligence**.

Step 2 Click the appropriate option:

- **Network Lists and Feeds** (for IP addresses)
- **DNS Lists and Feeds** (for domain names)
- **URL Lists and Feeds**

Step 3 Click the pencil beside the Global Block or Global Do-Not-Block list.

Step 4 Click the trash button beside the entry to delete.

List and feed updates for security intelligence feeds

A list and feed update is a security intelligence mechanism that performs the following functions:

- Replaces the existing list or feed file with the contents of the new file.
- Does not merge contents of existing and new files.
- If a corrupt or unrecognizable feed is downloaded, the system keeps using previous feed data. This is not true for the first download. If the system recognizes at least one entry in the feed, it uses only the entries it can recognize.

Feed update frequency and device polling

- Each feed updates the Cloud-Delivered Firewall Management Center every two hours by default; you can modify this frequency.
- When the Cloud-Delivered Firewall Management Center receives updates, it immediately passes them to your managed devices.
- Your managed devices poll the Cloud-Delivered Firewall Management Center every 30 minutes for changes. You cannot change this polling frequency.
- For Secure Firewall 200, updates are downloaded once a day, not every 30 minutes.

To modify feed update intervals, see Change the update frequency for security intelligence feeds, on page 3026.

Change the update frequency for security intelligence feeds

Set the intervals at which Security Intelligence feeds are updated in the system. Adjusting the update frequency ensures that threat data is refreshed according to your organization's needs.

You can specify the intervals at which the Cloud-Delivered Firewall Management Center updates Security Intelligence Feeds.

For details about feed updates, see List and feed updates for security intelligence feeds, on page 3025.

Before you begin

No explicit prerequisites are stated. Ensure you have access to the system interface where Security Intelligence feeds are managed.

Follow these steps to change the update frequency for Security Intelligence feeds:

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose the feed type whose frequency you want to change.

The system-provided URL feed is combined with the domain feed under **DNS Lists and Feeds**.

Step 3 Next to the feed you want to update, click **Edit** ( ).

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Edit the **Update Frequency**.

Step 5 Click **Save**.

The Security Intelligence feed will update at the new interval you specified. The system will retrieve threat data according to the configured frequency.

Custom security intelligence lists and feeds

This reference describes custom Security Intelligence lists and feeds.

List custom lists and feeds requirements

Each custom list or feed must be a simple text file no larger than 500 MB. List files must have the `.txt` extension. Include one entry or comment per line—for example, one IP address, one URL, or one domain name per line.

List and feed formatting

Each custom list or feed must be a simple text file no larger than 500 MB. List files must have the `.txt` extension. Include one entry or comment per line—for example, one IP address, one URL, or one domain name per line.



Tip The number of entries is limited by the maximum file size. For example, a URL list with no comments and an average URL length of 100 characters (including Unicode representations and newlines) can support more than 5.24 million entries.

- In a DNS list entry, you can use an asterisk (`*`) wildcard character for a domain label. All labels match the wildcard. Example: `www.example.*` matches `www.example.com` and `www.example.co`.
- Comment lines must start with the pound (`#`) character.
- Comments are removed during upload. Files you download contain only the entries, not the comments.

Feed requirements

- When configuring a feed, specify its location using a URL; the URL cannot be Punycode-encoded.
- For feed update intervals of 30 minutes or less, you must provide an MD5 URL to prevent frequent downloads of unchanged feeds. If your feed server does not provide an MD5 URL, use a download interval of at least 30 minutes.
- If you use an MD5 checksum, store it in a simple text file. The file should contain only the checksum. Do not add comments to the file.

List URL syntax and matching criteria for URL lists and feeds

Security Intelligence URL lists and feeds, including custom lists and feeds and entries in the global Block list and Do Not Block list, can include hostnames, URLs, slashes for exact matches, wildcards, and IP addresses, each with specific matching behavior.

- Hostnames

  For example, **`www.example.com`**.

- URLs
  - **`example.com`** matches **`example.com`** and all subdomains, including:
    - **`www.example.com`**
    - **`eu.example.com`**
    - **`example.com/abc`**
    - **`www.example.com/def`**
  - Does not match:
    - **`example.co.uk`**
    - **`examplexyz.com`**

- **example.com.malicious-site.com**

- When a feed or list includes a single entry, every URL that ends with those domains is identified and blocked
- Example feed:

  If you add `www.netflix.com, www.amazon.*, org, edu, www.hulu.*` to the feed, these are blocked:

  - `http://www.amazon.in`
  - `http://www.rajiv.org/`
  - `http://www.edu.edu/`
  - `http://www.edu.org/`
  - `http://org.org/`
  - `http://edu.edu`

- You can also include an entire URL path, such as **https://www.cisco.com/c/en/us/products/security/firewalls/index.html**

- Custom feeds with embedded credentials:

  You can create a custom URL, Network, and DNS feeds, including the username and password inside the URL (for example: **https://admin:password@server.domain.com/list.txt**).

Note If your password contains special characters (such as a colon (:) or an at sign (@)), transmission would fail. Ensure that your passwords do not contain these special characters, or use encoded passwords.

- A slash at the end of a URL to specify an exact match

  example.com/ matches only **example.com**; it does not match **www.example.com** or any other URL.

- A wildcard (*) to represent any domain in a URL
  - An asterisk can represent a complete domain string separated by dots, but not a partial domain string, and not any part of the URL following the first slash.

    Valid examples:

    - ***.example.com**
    - **www.*.com**
    - **example.***

      (This will match **example.com** and **example.org** and **example.de**, for example, but not **example.co.uk**)

    - ***.example.***

- **example.*/**

Invalid examples:

- **example*.com**
- **example.com/***

- IP addresses (IPv4)
  - For IPv6 addresses, or to use ranges or CIDR notation, use the Security Intelligence Network object.
  - You can include one or more wildcards representing an octet, for example 10.10.10.* or 10.10.*.*.

See also Custom Security Intelligence lists, on page 3031.

Custom security intelligence feeds

A custom security intelligence feed is a threat intelligence source that

- enables augmentation of system-provided intelligence feeds with regularly updated reputable block lists and do not block lists from the internet,
- supports configuration of internal feeds for updating multiple Cloud-Delivered Firewall Management Center appliances using a single source list, and
- allows use of MD5 checksums to determine if a feed requires downloading, reducing unnecessary updates for large or internal feeds.

Security intelligence feed management

Custom or third-party Security Intelligence feeds allow you to augment the system-provided intelligence feeds with other regularly updated reputable block lists and do not block lists on the internet. You can also set up an internal feed, which is useful if you want to update multiple Cloud-Delivered Firewall Management Center appliances in your deployment using one source list.

You can configure the system to use an MD5 checksum to determine whether to download an updated feed. If the checksum has not changed since the last download, the system does not need to re-download it. Using MD5 checksums is especially helpful for internal feeds that are large.

Note You cannot add address blocks to block or do not block lists using a `/0` netmask in a Security Intelligence feed. If you want to monitor or block all traffic targeted by a policy, use an access control rule with the **Monitor** or **Block** rule action and specify `any` for the **Source Networks** and **Destination Networks**.

Note The system does **not** perform peer SSL certificate verification when downloading custom feeds, nor does the system support the use of certificate bundles or self-signed certificates to verify the remote peer.

If you want strict control over when the system updates a feed from the internet, you can disable automatic updates for that feed. However, automatic updates ensure that the data is always current and relevant.

When you manually update Security Intelligence feeds, the system updates all feeds, including the system-provided intelligence feeds.

See complete requirements at List custom lists and feeds requirements, on page 3026.

Create custom security intelligence feeds

Configure security intelligence feeds to enhance your device's ability to detect and respond to threats by importing external or custom threat data sources.

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Use this task when you need to add new threat intelligence feeds for IP, DNS, or URL data to your security device.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Security Intelligence** node, then choose a feed type you want to add.

Step 2 Click the option appropriate to the feed type you chose:

- **Add Network Lists and Feeds** (for IP addresses)
- **Add DNS Lists and Feeds**
- **Add URL Lists and Feeds**

Step 3 Enter a **Name** for the feed.

Step 4 Choose **Feed** from the **Type** drop-down list.

Step 5 Enter a **Feed URL**.

Step 6 Enter an **MD5 URL**.

- The MD5 URL is used to check if the feed content has changed since the last update and is required for update intervals shorter than 30 minutes.
- If your feed provider does not supply an MD5 URL, set the update interval to at least 30 minutes.

Step 7 Choose an **Update Frequency**.

Step 8 Click **Save**.
Unless you disabled feed updates, the system downloads and verifies the feed.

Update custom security intelligence feeds manually

Manually updating Security Intelligence feeds allows you to ensure that your deployment uses the most current threat intelligence data for filtering network traffic.

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Before you begin

At least one device must already be added to the management center.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Security Intelligence**.

Step 2 Click **Update Feeds**, then confirm.

Step 3 Click **OK**.

After the Cloud-Delivered Firewall Management Center downloads and verifies the feed updates, it communicates any changes to its managed devices. Your deployment begins filtering traffic using the updated feeds.

Custom Security Intelligence lists

A custom Security Intelligence list is a static list that

- contains IP addresses, address blocks, URLs, or domain names that you manually upload to the system,
- enables you to augment and fine-tune feeds or global lists for the managed devices of a single Cloud-Delivered Firewall Management Center, and
- is useful when you need to override or supplement existing feed behavior without removing the entire feed from a policy.

Formatting requirements for custom Security Intelligence lists

- Netmasks for address blocks must be integers from `0` to `32` for IPv4 or from `0` to `128` for IPv6.
- Encode Unicode in domain names using Punycode format. The encoding is not case sensitive.
- Characters in domain names are not case sensitive.
- Unicode in URLs must be percent-encoded.
- Characters in URL subdirectories are case sensitive.
- List entries that start with the number sign (`#`) are treated as comments.
- For more information about formatting requirements, see the List custom lists and feeds requirements, on page 3026.

Matching requirements for custom Security Intelligence lists

- If you add a higher-level domain to a URL or DNS list, any sub-level domains also match. For example, if you add `example.com` to a DNS list, `www.example.com` and `test.example.com` also match.
- DNS lookups (forward or reverse) are not performed on DNS or URL list entries. For example, if you add `http://198.51.100.2` to a URL list, the list matches only that exact address and not any domain it resolves to.

Note You cannot add address blocks to a Block or Do Not Block list using a `/0` netmask in a Security Intelligence list. To monitor or block all traffic targeted by a policy, use an access control rule with the **Monitor** or **Block** action, and set the **Source Networks** and **Destination Networks** to `any`.

Examples for customizing Security Intelligence lists

- If a global feed blocks access to a vital resource but is otherwise useful, you can upload a custom list to allow access to that resource without disabling the feed entirely.
- You can supplement an existing feed by adding URLs or domains relevant to your organization.

Upload new custom security intelligence lists to the Cloud-Delivered Firewall Management Center

To modify a Security Intelligence list, you must make your changes to the source file and upload a new copy. You cannot modify the file's contents using the web interface. If you do not have access to the source file, download a copy from the system.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Security Intelligence** node, then choose a list type.

Step 2 Click the option appropriate to the list you chose above:

- **Add Network Lists and Feeds** (for IP addresses)
- **Add DNS Lists and Feeds**
- **Add URL Lists and Feeds**

Step 3 Enter a **Name**.

Step 4 From the **Type** drop-down list, choose **List**.

Step 5 Click **Browse** to browse to the list `.txt` file, then click **Upload**.

Step 6 Click **Save**.

What to do next

You do not need to redeploy these changes to take effect. If you want to delete an entry from the list, see Delete entries from global security intelligence lists, on page 3025.

Update custom security intelligence lists

Modify, add, or remove entries from your security intelligence lists in Secure Firewall Management Center. Update your lists when new threats emerge or when you need to refine your firewall's response to known risks. Perform this task as part of routine security maintenance or to update your firewall based on new threat intelligence.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Security Intelligence** node, then choose a list type.

Step 2 Next to the list you want to update, click **Edit** ( ).

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 To edit the list, click **Download** and save the list as a text file using your browser's prompts.

Step 4 Change the list as needed.

Step 5 On the Security Intelligence pop-up window, click **Browse** to browse to the modified list, then click **Upload**.

Step 6 Click **Save**.

What to do next

These changes take effect without redeployment. If you want to delete an entry from the list, see Delete entries from global security intelligence lists, on page 3025.

Sinkhole objects

A sinkhole object is a network security object that

- represents either a DNS server that gives non-routable addresses for all domain names within the sinkhole
- can represent an IP address that does not resolve to a server, and
- requires assignment of both an IPv4 address and an IPv6 address.

Sinkhole object usage in DNS policy rules

You can reference the sinkhole object within a DNS policy rule to redirect matching traffic to the sinkhole.

Create a sinkhole object

Create a sinkhole object to redirect or block suspicious traffic for enhanced threat detection and analysis.

Use this task when you need to configure a sinkhole object for traffic redirection or blocking to maintain security.

Before you begin

Ensure you have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management** > **Sinkhole**.

Step 2 Click **Add Sinkhole**.

Step 3 Enter a **Name**.

Step 4 Enter the **IPv4 Address** and **IPv6 Address** of your sinkhole.

Step 5 You have these options:

- To redirect traffic to a sinkhole server, choose **Log Connections to Sinkhole**.
- To redirect traffic to a non-resolving IP address, choose **Block and Log Connections to Sinkhole**.

Step 6 (Optional) From the **Type** drop-down list, assign an Indication of Compromise (IoC) type for the sinkhole object, if needed.

Step 7 Click **Save**.

Configure an SLA monitor

Configure an SLA monitor to define a connectivity policy to a monitored address and track the availability of a route. This helps ensure that the route is periodically checked and replaced with a backup route if it becomes unavailable.

An Internet Protocol Service Level Agreement (SLA) monitor defines a connectivity policy to a monitored address and tracks the availability of a route. The monitor periodically checks the route by sending ICMP echo requests. If requests time out, the route is removed from the routing table and replaced with a backup route. SLA monitoring jobs start immediately after deployment and continue to run unless you remove the SLA monitor from the device configuration. The SLA Monitor Object is used in the Route Tracking field of an IPv4 Static Route Policy.

Note IPv6 routes do not support SLA monitors via route tracking.

Procedure

Step 1 Select **Objects** > **Object Management** > **SLA Monitor**.

Step 2 Click **Add SLA Monitor** and enter these values:

Fields	Details
Name	Enter a name for the object.
Description	Enter a description for the object.

Fields	Details
Frequency	Enter the frequency of ICMP echo request transmissions, in seconds. Valid values range from 1 to 604800 seconds (7 days). The default is 60 seconds. **Note** The frequency cannot be less than the timeout value; you must convert frequency to milliseconds to compare the values.
SLA Monitor ID	Enter the ID number of the SLA operation. Values range from 1 to 2147483647. You can create a maximum of 2000 SLA operations on a device. Each ID number must be unique to the policy and the device configuration.
Threshold	Enter the amount of time that must pass after an ICMP echo request before a rising threshold is declared, in milliseconds. Valid values range from 0 to 2147483647 milliseconds. The default is 5000 milliseconds. The threshold value identifies events that exceed the defined value, which you can use to evaluate the proper timeout value. However, the threshold does not serve as a direct indicator of the reachability of the monitored address. **Note** The threshold value should not exceed the timeout value.
Timeout	Enter the amount of time that the SLA operation waits for a response to the ICMP echo requests, in milliseconds. Values range from 0 to 604800000 milliseconds (7 days). The default is 5000 milliseconds. If a response is not received from the monitored address within the amount of time defined in this field, the static route is removed from the routing table and replaced by the backup route. **Note** The timeout value cannot exceed the frequency value (adjust the frequency value to milliseconds to compare the numbers).
Data Size	Enter the size of the ICMP request packet payload, in bytes. Values range from 0 to 16384 bytes. The default is 28 bytes, which creates a total ICMP packet of 64 bytes. Do not set this value higher than the maximum allowed by the protocol or the Path Maximum Transmission Unit (PMTU). For purposes of reachability, you might need to increase the default data size to detect PMTU changes between the source and the target. A low PMTU can affect session performance and, if detected, might indicate that the secondary path should be used.
ToS	Enter a value for type of service (ToS) defined in the IP header of the ICMP request packet. Values range from 0 to 255. The default is 0. This field contains information such as delay, precedence, reliability, and so on. It can be used by other devices on the network for policy routing and features such as committed access rate.

Fields	Details
Number of Packets	Enter the number of packets that are sent. Values range from 1 to 100. The default is 1 packet. **Note** Increase the default number of packets if you are concerned that packet loss might falsely cause the Secure Firewall Threat Defense device to believe that the monitored address cannot be reached.
Monitored Address	Enter the IP address that is being monitored for availability by the SLA operation.

Step 3 The **Available Zones** list displays both zones and interface groups. In the **Zones/Interfaces** list, add the zones or interface groups that contain the interfaces through which the device communicates with the management station. To specify a single interface, you need to create a zone or the interface groups for the interface; see Create security zone and interface group objects, on page 1002. The host will be configured on a device only if the device includes the selected interfaces or zones.

Step 4 Click **Save**.

After completing these steps, the SLA monitor is configured and begins tracking the availability of the specified route. If the monitored address becomes unreachable, the route is removed from the routing table and replaced with a backup route.

Time ranges

A time range is a configuration object that

- defines specific time periods,
- enables use of these periods to determine when rules apply, and
- supports integration with time-based ACLs.

Time range objects allow you to specify periods for rule enforcement.

Note Time-based ACLs is supported in Snort 3 also from Cloud-Delivered Firewall Management Center 7.0 onwards.

Create a time range object

Establish a time range object so policies apply only during specified periods, enabling precise control of network traffic.

If you want a policy to apply only during a specified time range, create a time range object and specify that object in the policy. This object works on Firewall Threat Defense devices only.

You can specify time range objects only in certain policy types, which are listed in this topic.

Note The timezone represents the device's local time and is used only for applying the time ranges in rules in the policies that support the time ranges. The timezone does not change the configured time of the device. To verify the configuration, in the Firewall Threat Defense CLI, use the **show time-range timezone** and **show time** commands (see the *Cisco Secure Firewall Threat Defense Command Reference* guide). In addition, the timezone of a chassis overrides the management center timezone.

Before you begin

Time ranges operate based on the time zone associated with the device that processes the traffic. By default, this is UTC. To change the time zone associated with a device, go to **Devices** > **Platform Settings**.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Time Range** from the list of object types.

Step 2 Click **Add Time Range**.

Step 3 Enter the object name and specify the desired time periods.

- If you see a red error box around the object name you have entered, mouse over the **Name** field to see naming restrictions.
- All times are in UTC, unless you specify a time zone for the device in **Devices** > **Platform Settings**.
- Enter times using a 24-hour clock; for example, enter 13:30 for 1:30 PM.
- To specify a single continuous range, such as typical weekend hours (Fridays at 5pm through Mondays at 8am, including evenings and nights), choose Range Type **Range**.
- To specify part of multiple days, such as Monday through Friday from 8am to 5pm (excluding evenings, nights, and early mornings every day), choose Range Type **Daily Interval**.
- You may specify up to 28 periods per object.
- For multiple noncontiguous times of day or different hours for different days, create multiple recurring intervals. For example, to apply a policy at all times other than standard working hours, create a single time range object with the following two recurring intervals:
  - A Daily Interval for Monday through Friday from 5pm through 8am, and
  - A Range recurring interval for Friday at 5pm through Monday at 8am.

Step 4 Click **Save**.

What to do next

Configure time ranges in any of these:

- Access control rules
- Prefilter rules

- Tunnel rules
- VPN group policy

In a VPN group policy object, specify the time range object using the **Access Hours** field. For details, see Configure a group policy object, on page 3062 and Advanced group policy option fields and values, on page 3068.

Time zones

A time zone is a configuration object that

- specifies the local time zone for a managed device
- is used only for applying time ranges in rules in policies that support time ranges, such as access control, prefilter, and VPN group policies), and
- defaults to UTC if not assigned to a device.

Device support and usage of time zones

- Time zone objects are supported only for Firewall Threat Defense devices.
- Time-based ACLs is supported in Snort 3 also from Cloud-Delivered Firewall Management Center 7.0 onwards.

Tunnel zones

A tunnel zone is a network object that

- represents certain types of plaintext, passthrough tunnels that you explicitly tag for special analysis,
- is not an interface object, and
- can be used as an interface constraint in some configurations.

For detailed information, see Using tunnel zones to apply access control at the tunnel level, on page 2534.

URLs

A URL object is a configuration element that

- defines a single URL or IP address for use within access control policies, event searches, and Security Intelligence configurations,
- can be used in multiple locations throughout the system's web interface to specify what traffic is filtered or monitored, and
- provides distinct matching behavior compared to URL group objects, which define multiple URLs or addresses.

URL object configuration guidelines

Important For best practices for using this and similar options in Security Intelligence configurations and for URL rules in access control and QoS policies, see Manual URL filtering options, on page 2102.

When creating URL objects, consider the following guidelines:

- If the URL does not include a path (no / characters), the match is based only on the server's hostname. If a path is included (one or more / characters), the entire URL string is used for a substring match.
- A URL is considered a match if:
  - the string appears at the beginning of the URL,
  - the string follows a dot,
  - the string contains a dot at the beginning,
  - or the string follows the :// characters.
- The system ignores the encryption protocol (HTTP or HTTPS). Blocking a website applies to both unless an application condition targets a specific protocol. Specify URLs without the protocol (use example.com, not http://example.com).
- To match HTTPS traffic in access control rules, use the subject common name from the public key certificate. Do not include subdomain information. The subject common name in certificates may differ from the web site's domain (for example, youtube.com uses *.google.com). Use SSL Decryption policy to decrypt HTTPS traffic for consistent URL filtering results.
- URL objects will not match HTTPS traffic if the browser resumes a TLS session, as certificate information may not be available, which can lead to inconsistent results.
- When using URL objects as aliases in remote access VPN policies, ensure the URL is 128 characters or less.
- Do not use manual URL filtering to block or allow individual web pages or parts of sites (that is, URL strings with / characters), as servers can be reorganized and pages moved to new paths.

URL object matching examples

For example, ign.com matches ign.com or www.ign.com, but not versign.com. Use example.com rather than http://example.com when specifying a URL object.

Create a URL object

Use this task when you need to define a new URL object in the Secure Firewall Management Center for use in access control or other security policies.

Procedure

Step 1 Choose **Objects** > **Object Management** > **URL** from the list of object types.

Step 2 Choose **Add Object** from the **Add URL** drop-down list.

Step 3 Enter a **Name**.

Step 4 (Optional) Enter a **Description**.

Step 5 Enter the **URL** or IP address.

Step 6 (Optional) Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.
- If you want to add override values to this object, expand the Override section and click **Add**; see Add object overrides, on page 2942.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Variable sets

A variable set is a configuration group that

- collects and manages variables used in intrusion rules and policies
- allows customization of predefined and user-defined variables, and
- enables linking of variable sets to intrusion policies for flexible network security management.

Variable set management and usage

Variables represent values commonly used in intrusion rules to identify source and destination IP addresses and PORTS. You can also use variables in intrusion policies to represent IP addresses in rule suppressions, adaptive profile updates, and dynamic rule states.



Tip Preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules.

You use variable sets to manage, customize, and group your variables. You can use the default variable set provided by the system or create your own custom sets. Within any set you can modify predefined default variables and add and modify user-defined variables.

Most of the shared object rules and standard text rules that the system provides use predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web SERVERS use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.

Rules are more effective when variables more accurately reflect your network environment. At a minimum, you should modify default variables in the default set. By ensuring that a variable such as `$HOME_NET` correctly defines your network and `$HTTP_SERVERS` includes all web SERVERS on your network, processing is optimized and all relevant systems are monitored for suspicious activity.

To use your variables, you link variable sets to intrusion policies associated with access control rules or with the default action of an access control policy. By default, the default variable set is linked to all intrusion policies used by access control policies.

Adding a variable to any set adds it to all sets; that is, each variable set is a collection of all variables currently configured on your system. Within any variable set, you can add user-defined variables and customize the value of any variable.

Initially, the system provides a single, default variable set comprised of predefined default values. Each variable in the default set is initially set to its default value, which for a predefined variable is the value set by the Talos Intelligence Group and provided in rule updates.

Although you can leave predefined default variables configured to their default values, Cisco recommends that you modify a subset of predefined variables.

You could work with variables only in the default set, but in many cases you can benefit most by adding one or more custom sets, configuring different variable values in different sets, and perhaps even adding new variables.

When using multiple sets, it is important to remember that the *current value* of any variable in the default set determines the *default value* of the variable in all other sets.

When you select **Variable Sets** on the Object Manager page, the object manager lists the default variable set and any custom sets you created.

On a freshly installed system, the default variable set is comprised only of the default variables predefined by Cisco.

Each variable set includes the default variables provided by the system and all custom variables you have added from any variable set. Note that you can edit the default set, but you cannot rename or delete the default set.



Caution Importing an access control or an intrusion policy overwrites existing default variables in the default variable set with the imported default variables. If your existing default variable set contains a custom variable not present in the imported default variable set, the unique variable is preserved.

Related Topics

Manage variables, on page 3054
Manage variable sets, on page 3053

Variable sets in intrusion policies

A variable set is a configuration group that

- links to intrusion policies used in an access control policy,
- provides variable values for enabled intrusion rules, and
- affects deployment status when modified.

Deployment status for variable sets in intrusion policies

By default, the system links the default variable set to all intrusion policies used in an access control policy. When you deploy an access control policy that uses an intrusion policy, intrusion rules that you have enabled in the intrusion policy use the variable values in the linked variable set.

When you modify a custom variable set used by an intrusion policy in an access control policy, the system reflects the status for that policy as out-of-date on the Access Control Policy page. You must re-deploy the access control policy to implement changes in your variable set. When you modify the default set, the system reflects the status of all access control policies that use intrusion policies as out-of-date, and you must re-deploy all access control policies to implement your changes.

Variables

A variable is a configuration element that

- can be provided by the system or created by users
- is categorized as default, customized, or advanced, and
- enables flexible management of network addresses and ports in rule creation.

Variable categories

Variables are grouped into three categories: default, customized, and advanced.

- **Default variables** are provided by the system. You cannot rename or delete a default variable, and you cannot change its default value. However, you can create a customized version of a default variable.
- **Customized variables** are variables you create. These variables can include:
  - *customized default variables*

    When you edit the value for a default variable, the system moves the variable from the Default Variables area to the Customized Variables area. Variable values in the default set determine the default values of variables in custom sets. If you customize a default variable in the default set, the system changes the default value of the variable in all other sets.
  - *user-defined variables*

    You can add and delete your own variables, customize their values within different variable sets, and reset customized variables to their default values. When you reset a user-defined variable, it remains in the Customized Variables area.

    User-defined variables can be one of these types:
    - *network* variables specify the IP addresses of hosts in your network traffic.

- *port* variables specify TCP or UDP ports in network traffic, including the value `any` for either type.

- **Advanced variables** are provided by the system under specific conditions. These variables have a very limited deployment.

Variable usage examples

For example, if you create custom standard text rules, you might also want to add your own user-defined variables to more accurately reflect your traffic or as shortcuts to simplify the rule creation process. Alternatively, if you create a rule that you want to inspect traffic in the "demilitarized zone" (or DMZ) only, you can create a variable named `$DMZ` whose value lists the server IP addresses that are exposed. You can then use the `$DMZ` variable in any rule written for this zone.

List predefined default variables

By default, the system provides a single default variable set, which is comprised of predefined default variables. The Talos Intelligence Group uses rule updates to provide new and updated intrusion rules and other intrusion policy elements, including default variables. Because many intrusion rules from the system use predefined default variables, set appropriate values for these variables. You can modify the values for these default variables in any or all variable sets, depending on how you use variable sets to identify traffic on your network.

Caution Importing an access control or an intrusion policy overwrites existing default variables in the default variable set with the imported default variables. If your existing default variable set contains a custom variable not present in the imported default variable set, the unique variable is preserved.

This table describes the variables provided by the system and indicates which variables you typically would modify. For assistance determining how to tailor variables to your network, contact Professional Services or Support.

Table 163: System-Provided variables

Variable Name	Description	Modify?
`$AIM_SERVERS`	Defines known AOL Instant Messenger (AIM) SERVERS, and is used in chat-based rules and rules that look for AIM exploits.	Not required.
`$DNS_SERVERS`	Defines Domain Name Service (DNS) SERVERS. If you create a rule that affects DNS SERVERS specifically, you can use the `$DNS_SERVERS` variable as a destination or source IP address.	Not required in curre
`$EXTERNAL_NET`	Defines the network that the system views as the unprotected network, and is used in many rules to define the EXTERNAL network.	Yes, you should adeq `$HOME_NET` and then `$HOME_NET` as the val `$EXTERNAL_NET`.
`$FILE_DATA_PORTS`	Defines non-encrypted ports used in intrusion rules that detect files in a network stream.	Not required.

Variable Name	Description	Modify?
$FTP_PORTS	Defines the ports of FTP SERVERS on your network, and is used for FTP server exploit rules.	Yes, if your FTP SERVE ports other than the defau (you can view the defaul the web interface).
$GTP_PORTS	Defines the data channel ports where the packet decoder extracts the payload inside a GTP (General Packet Radio Service [GPRS] Tunneling Protocol) PDU.	Not required.
$HOME_NET	Defines the network that the associated intrusion policy monitors, and is used in many rules to define the internal network.	Yes, to include the IP add your internal network.
$HTTP_PORTS	Defines the ports of web SERVERS on your network, and is used for web server exploit rules.	Yes, if your web SERVE ports other than the defau (you can view the defaul the web interface).
$HTTP_SERVERS	Defines the web SERVERS on your network. Used in web server exploit rules.	Yes, if you run HTTP SE
$ORACLE_PORTS	Defines Oracle database server ports on your network, and is used in rules that scan for attacks on Oracle databases.	Yes, if you run Oracle SE
$SHELLCODE_PORTS	Defines the ports you want the system to scan for shell code exploits, and is used in rules that detect exploits that use shell code.	Not required.
$SIP_PORTS	Defines the ports of SIP SERVERS on your network, and is used for SIP exploit rules.	Not required.
$SIP_SERVERS	Defines SIP SERVERS on your network, and is used in rules that address SIP-targeted exploits.	Yes, if you run SIP SERV should adequately define $HOME_NET and then inclu $HOME_NET as the value f $SIP_SERVERS.
$SMTP_SERVERS	Defines SMTP SERVERS on your network, and is used in rules that address exploits that target mail SERVERS.	Yes, if you run SMTP SE
$SNMP_SERVERS	Defines SNMP SERVERS on your network, and is used in rules that scan for attacks on SNMP SERVERS.	Yes, if you run SNMP SE
$SNORT_BPF	Identifies a legacy advanced variable that appears only when it existed on your system in a software release before Version 5.3.0 that you subsequently upgraded to Version 5.3.0 or greater.	No, you can only view or variable. You cannot edit recover it after deleting i
$SQL_SERVERS	Defines database SERVERS on your network, and is used in rules that address database-targeted exploits.	Yes, if you run SQL SER
$SSH_PORTS	Defines the ports of SSH SERVERS on your network, and is used for SSH server exploit rules.	Yes, if your SSH SERVE ports other than the defau (you can view the defaul the web interface).

Variable Name	Description	Modify?
`$SSH_SERVERS`	Defines SSH SERVERS on your network, and is used in rules that address SSH-targeted exploits.	Yes, if you run SSH S you should adequatel `$HOME_NET` and then i `$HOME_NET` as the val `$SSH_SERVERS`.
`$TELNET_SERVERS`	Defines known Telnet SERVERS on your network, and is used in rules that address Telnet server-targeted exploits.	Yes, if you run Telne
`$USER_CONF`	Provides a general tool that allows you to configure one or more features not otherwise available via the web interface. Conflicting or duplicate `$USER_CONF` configurations will halt the system.	No, only as instructe description or with th Support.

Network variables

A network variable is a configuration element that

- represents IP addresses for use in intrusion rules, intrusion policy rule suppressions, dynamic rule states, and adaptive profile updates
- differs from network objects and network object groups by being specific to intrusion policies and intrusion rules, and
- enables specification of host IP addresses in various configurations within an intrusion policy.

Network variable usage and configuration

Network variables can be used to specify the IP addresses of hosts on the network in the following configurations:

- Intrusion rules: The **Source IPs** and **Destination IPs** fields restrict packet inspection to traffic originating from or destined to specific IP addresses.
- Suppressions: The **Network** field in source or destination intrusion rule suppressions suppresses intrusion event notifications when a specific IP address or range triggers an intrusion rule or preprocessor.
- Dynamic rule states: The **Network** field in dynamic rule states identifies when too many matches for an intrusion rule or preprocessor occur within a given time period.
- adaptive profile updates: The**Networks** field identifies hosts where packet fragment and TCP stream reassembly is improved in passive deployments.

When variables are used in these fields, the variable set linked to an intrusion policy determines the variable values applied to network traffic handled by the related access control policy.

A variable can include any combination of the following network configurations:

- network variables, network objects, and network object groups selected from the available networks list,
- individual network objects added from the New or Edit Variable page, which can then be used in variables or other configurations,

- literal single IP addresses or address blocks; multiple literal IP addresses and address blocks may be included individually, supporting both IPv4 and IPv6 addresses.

  When specifying IPv6 addresses, any addressing convention defined in RFC 4291 may be used.

The default value for included networks in any variable is `any`, which indicates any IPv4 or IPv6 address. The default value for excluded networks is `none`, meaning no network. If you specify the address `::`, it means any IPv6 address when included, or no IPv6 addresses when excluded.

Adding networks to the excluded list negates those addresses and blocks; all addresses except the excluded ones can be matched.

For example, excluding `192.168.1.1` specifies any IP address other than 192.168.1.1. Excluding `2001:db8:ca2e::fa4c`specifies any IP address except 2001:db8:ca2e::fa4c.

Exclusions can be made with any combination of literal or available networks. For instance, excluding both `192.168.1.1` and `192.168.1.5` matches any IP address other than 192.168.1.1 or 192.168.1.5.

Keep these points in mind when you add or edit network variables:

- The value `any` cannot be logically excluded, as it would mean no address; for example, it is not possible to add a variable with the value `any` to the exclusion list.
- Network variables identify traffic for specific intrusion rule and intrusion policy features; preprocessor rules can trigger events regardless of the hosts defined by network variables.
- Excluded values must resolve to a subset of included values; for example, including the address block 192.168.5.0/24 and excluding 192.168.6.0/24 is not valid.

Port variables

A port variable is a network configuration element that

- represents TCP and UDP ports you can use in the **Source Port** and **Destination Port** header fields in intrusion rules,
- is specific to intrusion rules and differs from port objects and port object groups, and
- can be used to restrict packet inspection to packets originating from or destined to specific TCP or UDP ports.

Port variable configuration details

Port variables are linked to intrusion policies associated with access control rules or policies, determining the values for these variables in the network traffic where you deploy the access control policy.

When configuring a port variable, you can include:

- any combination of port variables and port objects that you select from the list of available ports (excluding port object groups, which cannot be added to variables),
- individual port objects that you add while creating or editing a variable, which can also be reused for other variables,
- and single literal port values as well as port ranges (ranges must use a dash, e.g., 80-90; colon-separated ranges are supported for backward compatibility but not for new variables).

  You can list multiple literal port values and ranges by adding each individually in any combination.

Port variable configuration has these important considerations:

- The default value for included ports in any variable is `any`, indicating all ports or port ranges. The default for excluded ports is `none`, indicating no ports.

Tip To create a variable with the value `any`, name and save the variable without adding a specific value.

- You cannot exclude the value `any` from a variable, as that would result in no ports being matched.
- Adding ports to an excluded list prevents matching on those specified ports or port ranges. The system matches any port except those excluded.
- Excluded values must be a subset of the included values. For example, if you include the port range 10-50, you cannot exclude port 60.

Advanced variables

An advanced variable is a configuration mechanism that

- enables you configure of features not available through the web interface
- includes the USER_CONF variable as your primary advanced variable, and
- allows you to enter custom instructions or lines up to the maximum character length.

User_CONF variable details

The USER_CONF variable provides a general tool that allows configuration of one or more features not otherwise available via the web interface.

Caution Do **not** use the advanced variable USER_CONF to configure an intrusion policy feature unless you are instructed to do so in the feature description or by Support. Conflicting or duplicate configurations will halt the system.

When editing USER_CONF, you can type up to 4096 characters on a single line. The line wraps automatically. You can include any number of valid instructions or lines until you reach either the 8192-character maximum length for a variable or a physical limit, such as disk space. Use the backslash (\) character as a line continuation after any complete argument in a command directive.

Resetting USER_CONF empties it.

Variable resets

A variable reset is a configuration action that

- returns a variable to its default value on the variable set new or edit variables page,
- applies different reset behaviors depending on whether the variable is in a default or custom set, and
- affects the value used in any intrusion policy where the variable set is linked.

Variable reset values

You can reset a variable to its default value on the variable set new or edit variables page. This table summarizes the basic principles of resetting variables.

Table 164: Variable reset values

Resetting this variable type...	In this set type...	Resets it to...
default	default	the rule update value
user-defined	default	`any`
default or user-defined	custom	the current default set value (modified or unmodified)

When you reset a variable in a custom set, you return it to the current value for that variable in the default set.

If you reset or change a variable in the default set, you update the default value in all custom sets. If the reset icon is unavailable, you cannot reset the variable because it has no customized value in the set. If you have not customized the variable in a custom set, changing it in the default set also updates any intrusion policy where you have linked the variable set.

Note Before you modify a variable in the default set, check how this change affects any intrusion policy that uses the variable in a linked custom set, especially if you have not customized the variable value in the custom set.

You can hover your pointer over the **Reset** icon in a variable set to see the reset value. If the customized value and the reset value match, you are in one of these situations:

- you are in the custom or default set where you added the variable with the value `any`.
- you are in the custom set where you added the variable with an explicit value and elected to use the configured value as the default value.

Variable sets

A variable set is a configuration grouping that

- allows variables to be managed collectively across multiple sets
- enables the addition of variables to all sets simultaneously, and
- provides options for using configured or default values when adding variables.

Variable set behavior

When you add a variable to a variable set, it is also added to all other sets. If you add a variable from a custom set, decide if you want to use the configured value as the customized value in the default set.

- **If you use the configured value** (for example, 192.168.0.0/16), the variable is added to the default set using the configured value as a customized value with a default value of `any`. The current value in the default set determines the default value in other sets. Therefore, the initial default value in other custom sets is the configured value (which in this example is 192.168.0.0/16).

- **If you do not use the configured value**, the variable is added to the default set using only the default value `any` and, consequently, the initial, default value in other custom sets is `any`.

Examples: Adding user-defined variables to default sets

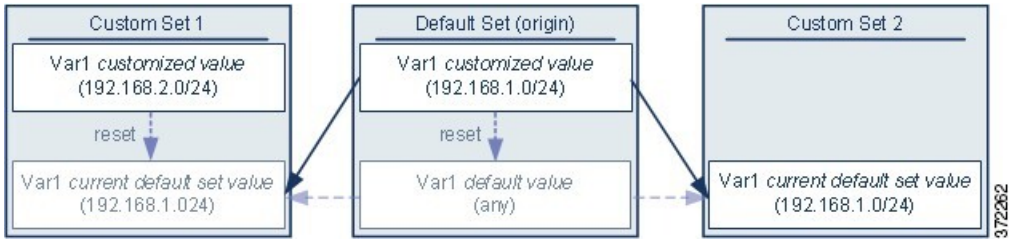
A user-defined variable addition is a configuration action that

- enables adding a variable, such as Var1, to a default set with a specific value (like 192.168.1.0/24),
- allows customization of the variable value in any set, and
- ensures that resetting the variable in the default set resets its default value to `any` across all sets.

Set interactions and variable customization

The diagram below illustrates how adding and customizing user-defined variables affects default and custom sets.

Figure 399: Set interaction diagram



Customizing the value of `Var1` in a custom set overrides the default value. For example, if `Var1` is customized in `Custom Set 1`, it takes on a new value, while in `Custom Set 2`, if `Var1` is not customized, it retains the default value (e.g., `192.168.1.0/24`). Resetting the user-defined variable in the default set updates its default value to `any` in all sets.

Further, when `Var1` in `Custom Set 2` is not updated, customizing or resetting `Var1` in the default set updates its current value in `Custom Set 2`, affecting any intrusion policy linked to the variable set.

Note While the interactions are the same for user-defined variables and default variables, resetting a default variable in the default set resets it to the value configured by Cisco in the current rule update, not `any`.

Customizing user-defined variable values

In Custom Set 1, the customized value `192.168.2.0/24` of `Var1` overrides the default value, while in Custom Set 2, `Var1` retains the default value `192.168.1.0/24` unless customized.

Resetting default variables in default sets

Although not shown in the example, note that interactions between sets are the same for user-defined variables and default variables except that resetting a default variable in the default set resets it to the value configured by Cisco in the current rule update.

Examples: User-defined variables in custom sets

A user-defined variable in a custom set is a configuration element that:

- enables you to add variables to custom sets,
- allows you to choose whether the configured value becomes the default for other sets, and
- affects how variable values and interactions propagate across sets.

Variable set interactions and default value behavior

When you add a user-defined variable to a custom set, you are prompted to decide if the configured value should be used as the default value for other sets. The examples illustrate the outcomes of each choice.

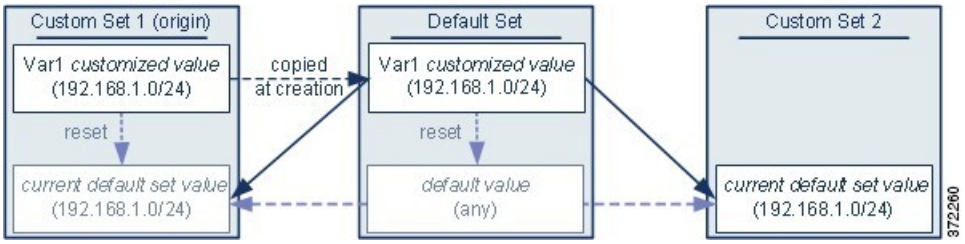
- Choosing to use the configured value as the default copies the value to the default set as a customized value, with a default value of `any`.
- Choosing not to use the configured value as the default adds the variable to all sets with a default value of `any`, allowing customization in any set.

Note Further customizing or resetting a variable in the default set updates the current, default value in other sets, which can affect any intrusion policy linked to the variable set.

Examples of adding user-defined variables to custom sets

The first example shows adding `Var1` with the value `192.168.1.0/24` to Custom Set 1 and electing **to use** the configured value as the default for other sets.

Figure 400: Variable set interaction when using the configured value as default

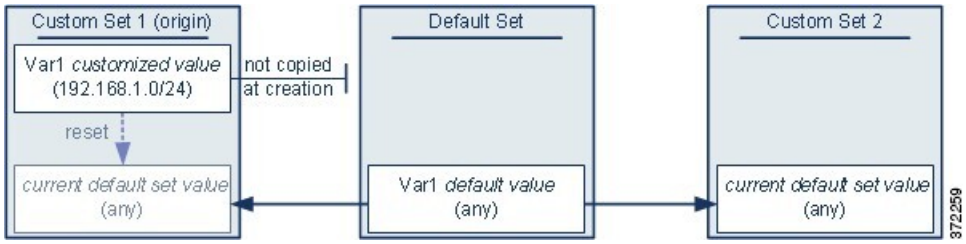


Except for the origin of `Var1` from Custom Set 1, this example is identical to adding `Var1` to the default set. Adding the customized value `192.168.1.0/24` for `Var1` to Custom Set 1 copies the value to the default set as a customized value with a default value of `any`. Thereafter, `Var1` values and interactions are the same as if you had added `Var1` to the default set. Further customizing or resetting `Var1` in the default set updates the current, default value of `Var1` in Custom Set 2, affecting any intrusion policy linked to the variable set.

Adding a variable and not using the configured value as default

The second example shows adding `Var1` with the value `192.168.1.0/24` to `Custom Set 1`, but electing **not to use** the configured value as the default in other sets.

Figure 401: Variable set interaction when not using the configured value as default



This approach adds `Var1` to all sets with a default value of `any`. After adding `Var1`, you can customize its value in any set. An advantage of this approach is that, by not initially customizing `Var1` in the default set, you decrease your risk of customizing the value in the default set and inadvertently changing the current value in a set such as Custom Set 2 where you have not customized `Var1`.

Nesting variables

A nesting variable is a configuration element that

- can include other variables as part of its definition
- does not support circular nesting, and
- does not support nested, negated variables.

Valid nested variables

In this example, SMTP_SERVERS, HTTP_SERVERS, and OTHER_SERVERS are valid nested variables.

Variable	Type	Included Networks	Excluded Networks
SMTP_SERVERS	customized default	10.1.1.1	—
HTTP_SERVERS	customized default	10.1.1.2	—
OTHER_SERVERS	user-defined	10.2.2.0/24	—
HOME_NET	customized default	10.1.1.0/24 OTHER_SERVERS	SMTP_SERVERS HTTP_SERVERS

Invalid nested variable

In this example, HOME_NET is an invalid nested variable because the nesting of HOME_NET is circular; that is, the definition of OTHER_SERVERS includes HOME_NET, so you would be nesting HOME_NET in itself.

Variable	Type	Included Networks	Excluded Networks
SMTP_SERVERS	customized default	10.1.1.1	—

Variable	Type	Included Networks	Excluded Networks
HTTP_SERVERS	customized default	10.1.1.2	—
OTHER_SERVERS	user-defined	10.2.2.0/24 HOME_NET	—
HOME_NET	customized default	10.1.1.0/24 OTHER_SERVERS	SMTP_SERVERS HTTP_SERVERS

Unsupported nested, negated variable

Because nested, negated variables are not supported, do not use the variable NONCORE_NET to represent IP addresses outside your protected networks.

Variable	Type	Included Networks	Excluded Networks
HOME_NET	customized default	10.1.0.0/16 10.2.0.0/16 10.3.0.0/16	—
EXTERNAL_NET	customized default	—	HOME_NET
DMZ_NET	user-defined	10.4.0.0/16	—
NOT_DMZ_NET	user-defined	—	DMZ_NET
NONCORE_NET	user-defined	EXTERNAL_NET NOT_DMZ_NET	—

Alternative to unsupported nested, negated variable

As an alternative to the example above, you could represent IP addresses that are outside of your protected networks by creating the variable NONCORE_NET to represent IP addresses outside your protected networks.

Variable	Type	Included Networks	Excluded Networks
HOME_NET	customized default	10.1.0.0/16 10.2.0.0/16 10.3.0.0/16	—
DMZ_NET	user-defined	10.4.0.0/16	—
NONCORE_NET	user-defined	—	HOME_NET DMZ_NET

Manage variable sets

Manage variable sets to organize and control the variables used by your devices. This task helps you add, edit, delete, or filter variable sets as needed for device configuration.

Variable sets are used to efficiently manage collections of variables crucial for device configuration. To use variable sets, you must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management** > **Variable Set**.

Step 2 Manage your variable sets:

- Add — If you want to add a custom variable set, click **Add Variable Set**; see Create a variable set, on page 3053.
- Delete — If you want to delete a custom variable set, click **Delete** ( ) next to the variable set, then click **Yes**. You cannot delete the default variable set or variable sets belonging to ancestor domains.

  Note
  Variables created in a variable set you delete are not deleted or otherwise affected in other sets.

- Edit — If you want to edit a variable set, click **Edit** ( ) next to the variable set you want to modify; see Edit objects, on page 2935.
- Filter — If you want to filter variable sets by name, begin entering a name; as you type, the page refreshes to display matching names. If you want to clear name filtering, click **Clear** ( ) in the filter field.
- Manage Variables — To manage the variables included in variable sets, see Manage variables, on page 3054.

Create a variable set

Variable sets allow you to define and group variables for use in Secure Firewall Management Center. Use this task when you need to create a new set of variables for policy configuration or reuse.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Variable Set**.

Step 2 Click **Add Variable Set**.

Step 3 Enter a **Name**.

Step 4 (Optional) Enter a **Description**.

Step 5 Manage the variables in the set; see Manage variables, on page 3054.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

Manage variables

Configure variable sets to control device policies and manage variables for Secure Firewall Management Center.

- You can display, add, delete, edit, or reset variables as required for your environment.

To manage variables in Secure Firewall Management Center, you must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management** > **Variable Set**.

Step 2 Click **Edit** ( ) next to the variable set you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Manage your variables:

- Display — If you want to display the complete value for a variable, hover your pointer over the value in the **Value** column next to the variable.
- Add — If you want to add a variable, click **Add**; see Add variables, on page 3055.
- Delete — Click **Delete** ( ) next to the variable. If you have saved the variable set since adding the variable, click **Yes** to confirm that you want to delete the variable.

  You *cannot* delete these:

  - default variables
  - user-defined variables that are used by intrusion rules or other variables
  - variables belonging to ancestor domains

- Edit — Click **Edit** ( ) next to the variable you want to edit; see Edit variables values in a variable set, on page 3056.
- Reset — If you want to reset a modified variable to its default value, click **Reset** next to a modified variable. If reset is dimmed, one of these is true:

  - The current value is already the default value.
  - The configuration belongs to an ancestor domain.

  Tip
  Hover your pointer over an active reset to display the default value.

Step 4 Click **Save** to save the variable set. If the variable set is in use by an access control policy, click **Yes** to confirm that you want to save your changes.

The current value in the default set determines the default value in all other sets. If you modify or reset a variable in the default set, the current value changes in other sets where you have not customized the default value.

What to do next

- If an active policy references your object, deploy configuration changes.

Add variables

Add variables to a variable set so you can define reusable network or port values for use in policies.

Use variables in device policies to enable flexible, reusable configurations that simplify network management.

Before you begin

Ensure you have the IPS license.

Procedure

Step 1 In the variable set editor, click **Add**.

Step 2 Enter a unique variable **Name**.

Step 3 From the **Type** drop-down list, choose either **Network** or **Port**.

Step 4 Define the values for the variable:

- To include or exclude specific networks or ports, select them and click **Include** or **Exclude**. If there is overlap, excluded addresses or ports always take precedence over included ones.
- To add literal values, enter a single IP address, address block, port, or port range. If you need to add more than one value, repeat this action.
- To remove any items from included or excluded lists, select the item and click **Delete** ( ) next to the item.
- You may use a combination of literal values, existing variables, objects, or network object groups when specifying included or excluded items for network variables.

Step 5 Save the variable.

- If adding a new variable from a custom set, choose whether to add it as a customized value in the default set or as a default value of `any` in the default and other custom sets.

Step 6 Save the variable set.

- Any access control policy linked to the variable set will display an out-of-date status until changes are deployed.

The variable is added to the variable set and is available for use in policies. Any access control policy linked to the variable set displays an out-of-date status until you deploy the changes.

What to do next

- If an active policy references your object, deploy configuration changes.

Edit variables values in a variable set

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

You can edit both custom and default variables.

You cannot change the **Name** or **Type** values in an existing variable.

Procedure

Step 1 In the variable set editor, click **Edit** ( ) next to the variable you want to modify.

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 2 Modify the variable:

- If you want to move items from the list of available networks or ports to the list of included or excluded items, you can select one or more items and then drag and drop, or click **Include** or **Exclude**.

  Tip
  If addresses or ports overlap between the included and excluded lists for a network or port variable, the excluded addresses or ports take precedence.

- Enter a single literal value, then click **Add**. For network variables, you can enter a single IP address or address block. For port variables, you can add a single port or port range, separating the upper and lower values with a hyphen (-). Repeat this step as needed to enter multiple literal values.
- If you want to remove an item from the included or excluded lists, click **Delete** ( ) next to the item.

Note
The list of items to include or exclude can be comprised of any combination of literal strings and existing variables, objects, and network object groups in the case of network variables.

Step 3 Click **Save** to save the variable.

Step 4 Click **Save** to save the variable set. If the variable set is in use by an access control policy, click **Yes** to confirm that you want to save your changes. Your changes are saved, and any access control policy the variable set is linked to displays an out-of-date status.

What to do next

- If an active policy references your object, deploy configuration changes.

VLAN tags

A VLAN tag is a network identifier that

- represents a single VLAN or a range of VLAN tags,
- can be grouped with other VLAN tag objects to create tag groups, and

VLAN tag usage

You can use VLAN tag objects and groups in various places in the system's web interface, such as rules and event searches. For example, you could write an access control rule that applies only to a specific VLAN.

Create a VLAN tag object

Use this task when you need to add or manage VLAN tag objects in your Secure Firewall Management Center. VLAN tag objects are referenced in access control and other policies to match specific VLAN tags or ranges.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VLAN Tag**.

Step 2 Choose **Add Object** from the **Add VLAN Tag** drop-down list.

Step 3 Enter a **Name**.

Step 4 Enter a **Description**.

Step 5 Enter a value in the **VLAN Tag** field. Use a hyphen to specify a range of VLAN tags.

Step 6 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allow object overrides, on page 2941.
- If you want to add override values to this object, expand the Override section and click **Add**; see Add object overrides, on page 2942.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes.

VPN

You can use the following VPN objects on Firewall Threat Defense devices. To use these objects, you must have Admin privileges, and your Smart License account must satisfy export controls. You can configure these objects in leaf domains only.

Certificate map objects

A certificate map object is a configuration element that

- contains a named set of certificate matching rules,
- associates a received certificate with a remote access VPN connection profile, and
- matches rules in priority order, ending when the first rule results in a match.

Certificate map object reference information

Remote access VPN policies use certificate map objects to map certificates to connection profiles. If a received certificate matches the rules in a certificate map, the system associates the connection with the specified profile.

Both certificate map objects and connection profiles are components of a remote access VPN policy.

- The rules are matched in the order shown in the UI.
- Matching ends when the first rule within the certificate map object results in a match.

Navigation path for certificate map objects:

- **Objects** > **Object Management** > **VPN** > **Certificate Map**

Fields in certificate map objects:

- **Name**: Identifies this object so it can be referred to from other configurations, such as remote access VPN.
- **Mapping criteria**: Specifies the contents of the certificate to evaluate. If the certificate satisfies these rules, the user is mapped to the connection profile containing this object.
  - **Field**: Select the field for the matching rule according to the subject or issuer of the client certificate.

    If you set the **Field** to *Alternative Subject* or *Extended Key Usage*, the component remains fixed as *Whole Field*.
  - **Component**: Select the component of the client certificate to use for the matching rule.



Note **SER (Serial Number) component**: Ensure you specify the serial number for the subject field. The certificate map only matches with a serial number attribute in the subject name.

  - **Operator**: Select the operator for the matching rule:
    - Equals: The certificate component must match the entered value. If they do not match exactly, the connection is denied.
    - Contains: The certificate component must contain the entered value. If the component does not contain the value, the connection is denied.

- Does not equal: The certificate component cannot equal the entered value. For example, for a selected certificate component of country and an entered value of US, if the client country value equals US, then the connection is denied.
- Does not contain: The certificate component cannot contain the entered value. For example, for a selected certificate component of country and an entered value of US, if the client country value contains US, the connection is denied.

- **Value**: The value of the matching rule. The value entered is associated with the selected component and operator.

Related Topics

Configure Certificate Maps, on page 1856

Secure Client custom attributes objects

Custom attributes objects enable Secure Client (formerly AnyConnect) to configure features such as Per App VPN, allow or defer upgrade, and dynamic split tunneling. You define the attribute type, then assign one or more named values for this type. In Secure Firewall Management Center, you create custom attributes objects, add them to a group policy, and associate the policy with a remote access VPN to enable features for VPN clients.

Supported features using custom attributes objects

Firewall Threat Defense supports these features using custom attribute objects:

- **Per App VPN**—You can specify which applications may use the VPN tunnel. The system identifies these apps and tunnels only the allowed applications.
- **Allow or defer upgrade**—You can delay downloading Secure Client updates. When a client update becomes available, you can prompt users to upgrade or defer the update.
- **Dynamic Split Tunneling**—You can configure policies to selectively include or exclude IP addresses or networks from the VPN tunnel. Create a custom attribute and add it to a group policy to achieve this.

For step-by-step instructions to configure Secure Client custom attributes, see Configure custom attributes for Secure Client, on page 3059.

For details about the specific custom attributes to configure for a feature, see the *Cisco Secure Client (including AnyConnect) Administrator Guide* for the Secure Client release you are using.

Configure custom attributes for Secure Client

Use this task when you need to configure custom attributes for Secure Client clients, such as Per App VPN, update deferment, or dynamic split tunneling, in Secure Firewall Management Center.

This is relevant for environments where device management and application-specific VPN rules are required, and when deploying secure client features that depend on custom attribute objects.

Applicable when integrating MDM-enrolled devices and advanced VPN features.

Before you begin

Ensure that you have done these steps before adding a custom attribute object for Per App VPN:

- Per App VPN must be properly configured via MDM and each device must be enrolled to the MDM server.
- Create a base64 encoded string for each app using the Cisco Secure Client Enterprise Application Selector Tool.
  1. Download the Cisco Secure Client Enterprise Application Selector Tool from here.
  2. Open the Application Selection Tool and select the mobile platform from the drop-down menu located on the upper left.
  3. Add rule by entering friendly name and App ID; the other fields are optional.
  4. On the menu bar, click **Policy**. The encoded base64 rule appears.
  5. Select and copy the policy string, and save it to use later when you create the Secure Client Custom Attributes object.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VPN** > **Custom Attribute**.

Step 2 Click **Secure Client Custom Attribute**.

Step 3 Enter a **Name** and optionally a **Description** for the attribute.

Step 4 Select an attribute from the **Secure Client Attribute** drop-down list:

- **Per App VPN**—Select this option and specify the base64 encoded string in the **Attribute Value** box.
- **Allow Defer Update**—Select one of these options and specify the required information to allow or defer Secure Client update:
  - **Show the prompt until user takes action**—Display the prompt to the VPN user until the user chooses to allow or defer the VPN client update.
  - **Show the prompt until times out**—Choose this option to display the prompt for a specified duration and specify the duration in the **Timeout** box.
  - **Do not show the prompt and take automatic action**—Choose this option to automatically allow or defer the VPN update.
  - **Default Action**—Select the default action to be taken when the user does not respond, or when you want to configure an automatic action without the user's intervention. You can choose to update the Secure Client or postpone the update.
  - **Minimum Version**—Specify the minimum Secure Client version to be present on the client system to allow or defer the update.
- **Dynamic Split Tunneling**—Select this option to include or exclude IP addresses or networks from the VPN tunnel.
  - **Include domains**—Specify domain names that will be included in the remote access VPN tunnel.
  - **Exclude domains**—Specify domain names that will be excluded from the remote access VPN tunnel.

Step 5 Select the **Allow Overrides** check box to enable object overrides.

Step 6 Click **Save**.
The custom attributes object is added to the list.

What to do next

Associate the custom attributes with a group policy. See Add custom attributes to your group policy, on page 3061.

Add custom attributes to your group policy

Add custom attributes to a group policy to enable their use with remote access VPN connections. You can assign advanced VPN features to users based on their group policy membership.

To use Secure Client custom attributes with remote access VPN connections, you must associate them with the appropriate group policy. This enables you to offer advanced VPN customization to users based on their group policy membership.

Procedure

Step 1 Select **Objects** > **Object Management** > **VPN** > **Group Policy**.

Step 2 Add a new group policy or edit an existing group policy.

Step 3 Click **Secure Client** > **Custom Attributes**.

Step 4 Click **Add**.

Step 5 Select the **Secure Client Attribute**: Per App VPN, Allow Defer Update, or Dynamic Split Tunneling.

Step 6 Select a **Custom Attribute Object** from the list.

Note
Click Add (+) to create a new custom attribute object for the selected Secure Client attribute. You can also create a custom attribute object at **Objects** > **Object Management** > **VPN** > **Custom Attribute**. See Configure custom attributes for Secure Client, on page 3059.

Step 7 Click **Add** to save the attributes to the group policy, and then click **Save** to save the changes to the group policy.

Related Topics

Group policy options for Secure Client, on page 3065

Firewall Threat Defense group policy objects

A group policy is a set of attribute and value pairs, stored in a group policy object, that define the remote access VPN experience. For example, in the group policy object, you configure general attributes such as addresses, protocols, and connection settings.

Group policy object assignment and licensing

The group policy applied to a user is determined when the VPN tunnel is established. The RADIUS authorization server assigns the group policy, or it is obtained from the current connection profile.



Note There is no group policy attribute inheritance on the Firewall Threat Defense. A group policy object is used, in its entirety, for a user. The group policy object identified by the AAA server upon login is used, or, if that is not specified, the default group policy configured for the VPN connection is used. The provided default group policy can be set to your default values, but will only be used if it is assigned to a connection profile and no other group policy has been identified for the user.

To use group objects, you must have one of these Secure Client licenses associated with your Smart License account with Export-Controlled Features enabled:

- Secure Client VPN Only
- Secure Client Advantage
- Secure Client Premier

Configure a group policy object

Group policy objects are used to manage and enforce security and access policies for remote access VPN connections.

See Firewall Threat Defense group policy objects, on page 3061.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VPN** > **Group Policy**.

Previously configured policies are listed including the system default. Depending on your level of access, you may edit, view, or delete a group policy.

Step 2 Click **Add Group Policy** or choose a current policy to edit.

Step 3 Enter a **Name** and optionally a **Description** for this policy.

The name can be up to 64 characters. Spaces are allowed. The description can be up to 1,024 characters.

Step 4 Specify the **General** parameters for this Group Policy as described in Group policy general options, on page 3063.

Step 5 Specify the **Secure Client** parameters for this Group Policy as described in Group policy options for Secure Client, on page 3065.

Step 6 Specify the **Advanced** parameters for this Group Policy as described in Advanced group policy option fields and values, on page 3068.

Step 7 Click **Save**.
The new Group Policy is added to the list.

What to do next

Add the group policy object to a remote access VPN connection profile.

Group policy general options

This topic describes the general options available when configuring a group policy, including navigation, VPN protocol selection, IP address pool assignment, banner configuration, DNS/WINS settings and split tunneling options.

Navigation path

Objects > **Object Management** > **VPN** > **Group Policy**, click **Add Group Policy** or choose an existing policy to edit, then select the **General** tab.

VPN protocols fields

Specify the types of Remote Access VPN tunnels (**SSL** or **IPsec IKEv2**) you want to apply for this group policy.

IP address pools

- Specifies the IPv4 address assignment based on address pools designated for user groups in Remote Access VPN.
- You can assign IP addresses from specific pools for identified user groups using RADIUS/ISE for authorization.
- You can enforce policy for users or user groups in non-identity-aware systems by configuring the Group Policy as the RADIUS Authorization attribute (GroupPolicy/Class).

Example: Select a specific address pool for contractors and enforce policy for restricted access to the internal network using those addresses.

The Firewall Threat Defense device assigns IPv4 address pools to the clients in this order:

1. RADIUS attribute for IPv4Address Pool
2. RADIUS attribute for Group Policy
3. Address Pool in Group Policy mapped to a Connection Profile
4. IPv4Address Pool in Connection Profile

Limitations:

- IPv6 address pool is not supported.
- A maximum of six IPv4 address pools can be configured in a Group Policy.
- Deployment failures may occur when address pools in use are modified. Before making any changes to the address pools, you must log off all users.
- If you rename address pools or configure overlapping address pools, deployment may fail. To address this, first remove the old address pool, then deploy the changed address pool.

  Troubleshooting commands :

- `show ip local pool <address-pool-name>`
- `show vpn-sessiondb detail anyconnect`
- `vpn-sessiondb loggoff all noconfirm`

Banner fields

Specifies the banner text to present to users at login. You can enter up to 491 characters in the banner. By default, no value is configured. The IPsec VPN client supports full HTML for the banner, however, the Secure Client supports only partial HTML. To ensure that the banner displays properly to remote users, use the /n tag for IPsec clients and the <BR> tag for SSL clients.

DNS/WINS fields

Domain Naming System (DNS) and Windows Internet Naming System (WINS) servers. Used for Secure Client name resolution.

- **Primary DNS Server** and **Secondary DNS Server**—Choose or create a Network Object which defines the IPv4 or IPv6 addresses of the DNS servers you want this group to use.
- **Primary WINS Server** and **Secondary WINS Server**—Choose or create a Network Object containing the IP addresses of the WINS servers you want this group to use.
- **DHCP Network Scope**—Choose or create a Network Object containing a routable IPv4 address on the same SUBNET as the desired pool, but not within the pool. The DHCP server determines which SUBNET this IP address belongs to and assigns an IP address from that pool. If not set properly, deployment of the VPN policy fails.

  If you configure DHCP servers for the address pool in the connection profile, the DHCP scope identifies the subnets to use for the pool for this group. The DHCP server must also have addresses in the same SUBNET identified by the scope. The scope allows you to select a subset of the address pools defined in the DHCP server for this specific group.

  If you do not define a network scope, the DHCP server assigns IP addresses in the order of the address pools configured. It goes through the pools until it identifies an unassigned address.

  We recommend using the IP address of an interface whenever possible for routing purposes. For example, if the pool is 10.100.10.2-10.100.10.254, and the interface address is 10.100.10.1/24, use 10.100.10.1 as the DHCP scope. Do not use the network number. Use DHCP only for IPv4 addressing. If the address you choose is not an interface address, you might need to create a static route for the scope address.

  LINK-SELECTION (RFC 3527) and SUBNET-SELECTION (RFC 3011) are currently not supported.
- **Default Domain**—Name of the default domain. Specify a top-level domain, for example, example.com.

Split tunneling fields

Split tunneling directs some network traffic through the VPN tunnel (encrypted) and the remaining network traffic outside the VPN tunnel (unencrypted or “in the clear”).

- **IPv4 Split Tunneling / IPv6 Split Tunneling**—By default, split tunneling is not enabled. For both IPv4 and IPv6 it is set to **Allow all traffic over tunnel**. Left as is, all traffic from the endpoint goes over the VPN connection.

To configure split tunneling, choose the **Tunnel networks specified below** or **Exclude networks specified below** policy, and then configure an access control list for that policy.

- **Split Tunnel Network List Type**—Choose the type of Access List you are using. Then choose or create a **Standard Access List** or **Extended Access List**. See Access lists, on page 2948 for details.

Note From Version 7.4, the **Standard Access List** option is disabled by default. Use the **Extended Access List** option to add an access list.

- **DNS Request Split Tunneling**—Also known as Split DNS. Configure the DNS behavior expected in your environment.

  By default, split DNS is not enabled and set to **Send DNS request as per split tunnel policy**. Choosing **Always send DNS request over tunnel** forces all DNS requests to be sent over the tunnel to the private network.

  To configure split DNS, choose **Send only specified domains over tunnel**, and enter the list of domain names in the **Domain List** field. These requests are resolved through the split tunnel to the private network. All other names are resolved using the public DNS server. Enter up to 10 entries in the list of domains, separated by commas. The entire string can be no longer than 255 characters.

Related Topics

Configure a group policy object, on page 3062

Group policy options for Secure Client

This topic lists and describes the specifications and configuration fields for group policy Secure Client VPN operation. This includes navigation steps, profile fields, management profile settings, client modules, SSL settings, connection options, and custom attributes.

These specifications apply to the operation of the Secure Client VPN.

Navigation

Objects > **Object Management** > **VPN** > **Group Policy**. Click **Add Group Policy** or choose a current policy to edit., and then select the **Secure Client** tab.

Profile fields

Profile—Choose or create a file object containing the Secure Client Profile. See File objects, on page 3075 for object creation details.

The Secure Client Profile is a group of configuration parameters stored in an XML file. The Secure Client software uses it to configure the connection entries that appear in the client's user interface. These parameters (XML tags) also configure settings to enable more Secure Client features.

Use the GUI-based Secure Client Profile Editor, an independent configuration tool, to create the Secure Client Profile. See the *Secure Client Profile Editor* chapter in the appropriate release of the Cisco Secure Client (including AnyConnect) Administrator Guide for details.

Management profile fields

Management VPN Tunnel—Provides always-on connectivity to the corporate network when the endpoint is powered up, regardless of the user's VPN connection status.

Management VPN Profile—The Management Profile file contains settings for enabling and establishing Management VPN Tunnel on endpoint.

The Standalone Management VPN Tunnel profile editor can be used to create a new profile file or modify an existing file. You can download the profile editor from Cisco Software Download Center.

For more information about adding a profile file, see File objects, on page 3075.

Client modules fields

Cisco Secure Client VPN Only offers enhanced security through various built-in modules. These modules provide services such as web security, network visibility into endpoint flows, and off-network roaming protection. Each module includes a custom client profile.

The following Secure Client modules are optional and you can configure these modules to be downloaded when a VPN user downloads Secure Client:

- **AMP Enabler**—Deploys advanced malware protection (AMP) for endpoints.
- **DART**—Captures a snapshot of system logs and other diagnostic information, which can be sent to the Cisco TAC for troubleshooting.
- **ISE Posture**—Uses the OPSWAT library to perform posture checks to assess an endpoint's compliance.
- **Network Access Manager**—Provides 802.1X (Layer 2) and device authentication for access to both wired and wireless networks.
- **Network Visibility**—Enhances the enterprise administrator's ability to do capacity and service planning, auditing, compliance, and security analytics.
- **Start Before Login**—Forces the user to connect to the enterprise infrastructure over a VPN connection before logging on to Windows by starting Secure Client before the Windows login dialog box appears.
- **Umbrella Roaming Security**—Provides DNS-layer security when no VPN is active.
- **Web Security**—Enforces web security policies and blocks malicious content.

Click **Add** and select the following for each client module:

- **Client Module**—Select the Secure Client module from the list.
- **Profile to download**—Choose or create a file object containing the Secure Client Profile. See File objects, on page 3075 for object creation details.
- **Enable module download**—Select to enable endpoints to download the client module along with the profile. If not selected, the endpoints can download only the client profile.

Use the GUI-based Secure Client Profile Editor, an independent configuration tool to create a client profile for each module. Download the Secure Client Profile Editor from Cisco Software Download Center. See the *Secure Client Profile Editor* chapter in the appropriate release of the Cisco Secure Client (including AnyConnect) for details.

SSL settings fields

- **SSL Compression**—Determines if data compression is enabled, and specifies the method to use (Deflate or LZS). SSL Compression is disabled by default.

  Data compression speeds up transmission rates, but also increases the memory requirement and CPU usage for each user session. Thereby, decreasing the overall throughput of the security appliance.

- **DTLS Compression**—Specifies whether to compress Datagram Transport Layer Security (DTLS) connections for this group using LZS. DTLS Compression is disabled by default.

- **MTU Size**—TSpecifies the maximum transmission unit (MTU) size for SSL VPN connections established by Cisco Secure Client VPN Only. The default is 1406 bytes; the valid range is 576 to 1462 bytes.

  - **Ignore DF Bit**—Whether to ignore the Don't Fragment (DF) bit in packets that need fragmentation. Allows the forced fragmentation of packets that have the DF bit set, allowing them to pass through the tunnel.

Connection settings fields

- **Enable Keepalive Messages between Secure Client and VPN gateway** and its **Interval** setting.—Determines whether peers exchange keepalive messages to demonstrate availability for sending and receiving data in the tunnel. By default, this feature is enabled. Keepalive messages are transmitted at set intervals. If enabled, enter the time interval (in seconds) that the remote client waits between sending IKE keepalive packets. The default interval is 20 seconds, the valid range is 15 to 600 seconds.

- **Enable Dead Peer Detection** and its **Interval** settings.—Dead Peer Detection (DPD) ensures that the VPN secure gateway or client quickly detects when the peer is no longer responding and the connection has failed. Default is enabled for both the gateway and the client. DPD messages transmit at set intervals. If enabled, enter the time interval (in seconds) that the remote client waits between sending DPD messages. The default interval is 30 seconds, the valid range is 5 to 3600 seconds.

- **Enable Client Bypass Protocol**—Allows you to configure how the secure gateway manages IPv4 traffic (when it is expecting only IPv6 traffic), or how it manages IPv6 traffic (when it is expecting only IPv4 traffic).

  After the Secure Clientconnects to the headend, the headend assigns the VPN connection an IPv4 address, an IPv6 address, or both. If the headend assigns only one type of address, you can configure the Client Bypass Protocol to either drop network traffic for which an IP address was not assigned (default, disabled, not checked), or allow that traffic to bypass the headend and be sent from the client unencrypted (“in the clear”; enabled, checked).

  For example, if the secure gateway assigns only an IPv4 address to the Secure Client connection and the endpoint is dual-stacked. When the endpoint attempts to reach an IPv6 address, if Client Bypass Protocol is disabled, the IPv6 traffic is dropped; however, if Client Bypass Protocol is enabled, the IPv6 traffic is sent from the client in the clear.

- **SSL rekey**—Enables the client to rekey the connection, renegotiating the crypto keys and initialization vectors, increasing the security of the connection. This is disabled by default. When enabled, the renegotiation can be done at a specified interval and rekey the existing tunnel or create a new tunnel by setting the following fields:

  - **Method**—Available when SSL rekey is enabled. Create a **New Tunnel** (default), or renegotiate, the **Existing Tunnel**'s specifications.

- **Interval**—Available when SSL rekey is enabled. Set to a default of 4 minutes with a range of 4-10080 minutes (1 week).

- **Client Firewall Rules**—Use the Client Firewall Rules to configure firewall settings for the VPN client's platform. Rules are based on criteria such as source address, destination address, and protocol. Extended Access Control List building block objects are used to define the traffic filter criteria. Choose or create an Extended ACL for this group policy. Define a **Private Network Rule** to control data flowing to the private network, a **Public Network Rule** to control data flowing "in the clear", outside of the established VPN tunnel, or both.

Note Ensure that the ACL contains only TCP/UDP/ICMP/IP ports and source network as any, any-ipv4 or any-ipv6.

Only VPN clients running Microsoft Windows can use these firewall settings.

Custom attributes fields

Custom attributes are used by the Secure Client to configure features such as Per App VPN, Allow or defer upgrade, and Dynamic split tunneling. Click **Add** to add custom attributes to the group policy.

1. Select the **Secure Client Attrinute**: Per App VPN, Allow Defer Update, or Dynamic Split Tunneling.
2. Select a **Custom Attribute Object** from the list.

Note Click Add (+) to create a new custom attribute object for the selected Secure Client attribute. You can also create a custom attribute object at **Objects** > **Object Management** > **VPN** > **Custom Attribute**. See Configure custom attributes for Secure Client, on page 3059.

3. Click **Add** to save the attributes to the group policy, and then click **Save** to save the changes to the group policy.

Related Topics

Configure a group policy object, on page 3062

Advanced group policy option fields and values

The advanced group policy includes options and attributes for configuring VPN access and session behavior. This reference describes where to locate advanced settings, outlines attributes related to traffic filtering, and summarizes session control fields and their allowed values.

Navigation path

Objects > **Object Management** > **VPN** > **Group Policy**. Click **Add Group Policy** or choose a current policy to edit. Then select the **Advanced** tab.

Traffic filter fields

- **Access List Filter**—Filters consist of rules that determine whether to allow or block tunneled data packets coming through the VPN connection. You set rules using criteria such as source address, destination address, and protocol.

  The VPN filter applies only to initial connections. It does not control secondary connections, such as a SIP media connection, that are opened during application inspection.

  Use Extended Access Control List building block objects to define traffic filter criteria. Choose or create an Extended ACL for this group policy.

- **Restrict VPN to VLAN**—Also called "VLAN mapping," this parameter specifies the egress VLAN interface for sessions to which this group policy applies. The ASA forwards all group traffic to the VLAN you select.

  Use this attribute to assign a VLAN to the group policy to simplify access control. You can assign a value to this attribute instead of using ACLs to filter session traffic. The drop-down list shows only the VLANs that are configured in this ASA. Allowed values range from 1 to 4094. The default value is Unrestricted.

Session settings fields

- **Access Hours**—Choose or create a time range object. This object specifies the range of time this group policy is available to be applied to a remote access user. See Time ranges, on page 3036 for details.
- **Simultaneous Logins Per User**—Specifies the maximum number of simultaneous logins allowed for a user. The default value is 3. If the minimum value is set to 0, logins are disabled and users cannot access the VPN. If you allow several simultaneous connections, security may be compromised and performance reduced.
- **Maximum Connection Time / Alert Interval**—Specifies the maximum user connection time in minutes. At the end of this time, the system disconnects the user. The minimum is 1 minute). The Alert interval lets you display a message to the user before the maximum connection time is reached.
- **Idle Timeout / Alert Interval**—Specifies this user's idle timeout period in minutes. If there is no communication activity on the user connection in this period, the system stops the connection. The minimum time is 1 minute. The default is 30 minutes. The Alert interval lets you display a message to the user before the idle time is reached.

Related Topics

Configure a group policy object, on page 3062

Firewall Threat Defense IPsec proposals

IPsec Proposals (or Transform Sets) are used when configuring VPN topologies. During the IPsec security association negotiation with ISAKMP, both peers must agree on the same proposal to protect the data flow.

IPsec proposal types and negotiation

IPsec proposals are categorized based on the IKE version: IKEv1 or IKEv2.

- For an IKEv1 IPsec proposal (Transform Set), select the mode for IPsec operation and define the required encryption and authentication types. Select only one option for each algorithm. To support multiple combinations in a VPN, create multiple IKEv1 IPsec proposal objects.

- For an IKEv2 IPsec proposal, select all encryption and hash algorithms allowed in the VPN. During IKEv2 negotiations, the peers negotiate and use the most appropriate options supported by both devices.

Both IKEv1 and IKEv2 IPsec proposals use the Encapsulating Security Protocol (ESP), which provides authentication, encryption, and antireplay services. ESP uses IP protocol type 50.

Note We recommend that you use both encryption and authentication on IPsec tunnels.

Configure an IKEv1 IPsec proposal object

Use this task when you need to create or modify IKEv1 IPsec proposal objects in Secure Firewall Management Center. Proposal objects define the cryptographic parameters for VPN connections and are required for establishing secure tunnels between firewalls or security gateways.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VPN** > **IKEv1 IPsec Proposal** and then **VPN** > **IPsec IKev1 Proposal** from the table of contents.

Previously configured Proposals are listed including system defined defaults. Depending on your level of access, you may **Edit** ( ), **View** ( ), or **Delete** ( ) a Proposal.

Step 2 Choose **Add IPsec IKEv1 Proposal** to create a new Proposal.

Step 3 Enter a **Name** for this proposal.

The name of the policy object can be up to 128 characters.

Step 4 Enter a **Description** for this Proposal.

Provide a description of the policy object that is up to 1024 characters.

Step 5 Choose the **ESP Encryption** method. The Encapsulating Security Protocol (ESP) encryption algorithm for this Proposal.

For IKEv1, select one of the options. When deciding which encryption and Hash Algorithms to use for the IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN. For a full explanation of the options, see Decide encryption algorithms for VPN policies, on page 1715.

Step 6 Select an option for **ESP Hash**.

For a full explanation of the options, see Decide which hash algorithms to use, on page 1715.

Step 7 Click **Save**
The new Proposal is added to the list.

Configure an IKEv2 IPsec proposal object

Use this task when you need to define or update the cryptographic settings for IKEv2 IPsec VPN tunnels on your Secure Firewall Management Center. This is relevant when setting up new VPNs or adjusting security policies to meet organizational requirements.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VPN** > **IKEv2 IPsec Proposal** from the table of contents.

Previously configured proposals are listed including system defined defaults. Depending on your level of access, you may **Edit** ( ), **View** ( ), or **Delete** ( ) a Proposal.

Step 2 Choose **Add IKEv2 IPsec Proposal** to create a new proposal.

Step 3 Enter a **Name** for this proposal.

The policy object name can be up to 128 characters.

Step 4 Enter a **Description** for this proposal.

The policy object description can contain up to 1,024 characters.

Step 5 Click **ESP Hash**, the hash or integrity algorithm to use in the proposal for authentication.

Note
Firewall Threat Defense does not support IPSec tunnels with NULL encryption. Make sure that you do not choose NULL encryption for IPSec IKEv2 proposal.

For IKEv2, select all the options you want to support for **ESP Hash**. For a full explanation of the options, see Decide which hash algorithms to use, on page 1715.

Step 6 Click **ESP Encryption** . Select the Encapsulating Security Protocol (ESP) encryption algorithm for this proposal.

For IKEv2, click Select to open a dialog box where you can select all of the options you want to support. When deciding which encryption and hash algorithms to use for the IPsec proposal, choose only algorithms supported by the devices in the VPN. For a full explanation of the options, see Decide encryption algorithms for VPN policies, on page 1715.

Step 7 Click **Save**.
The new proposal is added to the list.

IKE policies

A IKE policy is a key management policy that

- defines security parameters and algorithms used to authenticate IPsec peers,
- negotiates and distributes IPsec encryption keys, and
- establishes secure communications by negotiating security associations in IPsec VPNs.

- The Internet Key Exchange (IKE) protocol manages authentication, key negotiation, and automatic establishment of IPsec security associations (SAs).
- IKE negotiation occurs in two phases:
  - Phase 1 negotiates a security association between two IKE peers, enabling secure communication in Phase 2.
  - Phase 2 establishes SAs for applications, such as IPsec.
- Both phases use proposals to negotiate connections. An IKE proposal is a set of algorithms that two peers use to secure negotiation.
- IKE negotiation starts with peers agreeing on a shared IKE policy, which defines the security parameters to protect future negotiations.
- For IKEv1, proposals include a single set of algorithms and a modulus group; multiple, prioritized policies can be created to match remote peers.
- For IKEv2, you can select multiple algorithms and modulus groups in one policy.
- Peers choose proposals during Phase 1 negotiation, enabling creation of a single IKE proposal but consideration of multiple prioritized options.
- For IKEv2, the policy object does not specify authentication; other policies must define authentication requirements.
- An IKE policy is required when you configure a site-to-site IPsec VPN. .

Configure an IKEv1 policy object

Use the IKEv1 Policy page to create, edit, or delete IKEv1 policy objects. These policy objects specify settings such as encryption, authentication, priority, and other parameters necessary for establishing secure VPN tunnels.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VPN** > **IKEv1 Policy** from the table of contents.

Previously configured policies are listed including system-defined defaults. Depending on your level of access, you may **Edit** ( ), **View** ( ), or **Delete** ( ) a proposal.

Step 2 (Optional) Click **Add IKEv1 Policy** to create a new policy object. Then, enter a **Name** for this policy. The policy object name can be up to 128 characters. Optionally, you can enter the description. The policy object description can contain up to 1,024 characters.

Step 3 Enter the **Priority** value of the IKE policy.

The priority value determines the order of the IKE policy compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, it tries to use the parameters defined in the next lowest priority. Valid values range from 1 to 65,535. The lower the number, the higher the priority. If you leave this field blank, Management Center assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

Step 4 Select the **Encryption** method.

When deciding which encryption and Hash Algorithms to use for the IKEv1 policy, your choice is limited to algorithms supported by the peer devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. For IKEv1, select one of the options. For a full explanation of the options, see Decide encryption algorithms for VPN policies, on page 1715.

Step 5 Select the **Hash** Algorithm that creates a Message Digest, which is used to ensure message integrity.

When deciding which encryption and Hash Algorithms to use for the IKEv1 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. For a full explanation of the options, see Decide which hash algorithms to use, on page 1715.

Step 6 Select the **Diffie-Hellman Group**.

The Diffie-Hellman group to use for encryption. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select the group that you want to allow in the VPN. For a full explanation of the options, see Decide which Diffie-Hellman modulus group to use, on page 1716.

Step 7 Enter the **Lifetime** of the security association (SA), in seconds. You can specify a value from 120 to 2,147,483,647 seconds. The default is 86400.

When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. Generally, the shorter the lifetime (up to a point), the more secure your IKE negotiations. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes.

Step 8 Select the **Authentication Method** to use between the two peers.

- **Preshared Key**—Preshared keys allow for a secret key to be shared between two peers and to be used by IKE during the authentication phase. If one of the participating peers is not configured with the same preshared key, the IKE SA cannot be established.
- **Certificate**—When you use Certificates as the authentication method for VPN connections, peers obtain digital certificates from a CA server in your PKI infrastructure, and trade them to authenticate each other.

Note
In a VPN topology that supports IKEv1, the **Authentication Method** specified in the chosen IKEv1 Policy object becomes the default in the IKEv1 **Authentication Type** setting. These values must match, otherwise, your configuration will error.

Step 9 Click **Save**
The new IKEv1 policy is added to the list.

Configure an IKEv2 policy object

Use the IKEv2 policy dialog box to create, delete, and edit an IKEv2 policy object. These policy objects contain the parameters required for IKEv2 policies.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VPN** > **IKEv2 Policy** from the table of contents.

Previously configured policies are listed including system defined defaults. Depending on your level of access, you may **Edit** (), **View** (), or **Delete** ( ) a policy.

Step 2 Click **Add IKEv2 Policy** to create a new policy. Then, enter a **Name** for this policy. The policy object name can be up to 128 characters. Optionally, enter the description. The description can be up to 1024 characters.

Step 3 Enter the **Priority**.

The priority value of the IKE proposal. The priority value determines the order of the IKE proposals compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, it tries to use the parameters defined in the next lowest priority policy. Valid values range from 1 to 65535. The lower the number, the higher the priority. If you leave this field blank, Management Center assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

Step 4 Enter the **Lifetime** of the security association (SA), in seconds. You can specify a value from 120 to 2,147,483,647 seconds. The default is 86400.

When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. Generally, the shorter the lifetime (up to a point), the more secure your IKE negotiations. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes.

Step 5 Select algorithms for these methods:

a. Click **Integrity Algorithms** portion of the Hash Algorithm used in the IKE policy. The Hash Algorithm creates a Message Digest, which is used to ensure message integrity.

When deciding which encryption and Hash Algorithms to use for the IKEv2 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. Select all the algorithms that you want to allow in the VPN.For a full explanation of the options, see Decide which hash algorithms to use, on page 1715.

b. Click **Encryption Algorithm** used to establish the Phase 1 SA for protecting Phase 2 negotiations.

When deciding which encryption and Hash Algorithms to use for the IKEv2 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. Select all the algorithms that you want to allow in the VPN. For a full explanation of the options, see Decide encryption algorithms for VPN policies, on page 1715.

c. Click **PRF Algorithm**.

The pseudorandom function (PRF) portion of the Hash Algorithm used in the IKE policy. In IKEv1, the Integrity and PRF algorithms are not separated, but in IKEv2, you can specify different algorithms for these elements. Select all of the algorithms that you want to allow in the VPN. For a full explanation of the options, see Decide which hash algorithms to use, on page 1715.

d. Click **Diffie-Hellman Group** for encryption.

A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select the groups that you want to allow in the VPN. For a full explanation of the options, see Decide which Diffie-Hellman modulus group to use, on page 1716.

Step 6 Click **Save**.
If a valid combination of choices has been selected the new IKEv2 policy is added to the list. If not, errors are displayed and you must make changes accordingly to successfully save this policy.

Secure client customizations

A Secure Client customization is a configuration object that

- represents files used to customize Secure Client,
- enables deployment of customizations to the VPN headend, and
- allows threat defense to distribute customizations to endpoints when users connect from Secure Client.

Supported secure client customizations

Secure Client customization objects represent files used to modify Secure Client. Secure Client supports these types of customizations:

- GUI text and messages
- Icons and images
- Scripts
- Binaries
- Custom installer transforms
- Localized installer transforms

For more information about configuring Secure Client customizations, see Customize Cisco Secure Client, on page 1868.

File objects

File objects are configuration element that

- represent files used in configurations, typically for remote access VPN policies
- can contain Secure Client Profile and Secure Client Image files, and
- are created, edited, and managed through the Add and Edit File Object dialog boxes.

File object management and usage

File objects are used to manage files for remote access VPN policies and are associated with profiles and images for deployment to endpoints.

File objects can be created and edited using the Add and Edit File Object dialog boxes. They are used in configurations for remote access VPN policies and can contain Secure Client Profile and Secure Client Image files.

Create profiles for each AnyConnect module and Secure Client Management VPN using independent profile editors. Deploy these profiles to meet administrator-defined end user requirements and authentication policies on endpoints Secure Client only. This makes preconfigured network profiles available to end users

When you create a file object, the Cloud-Delivered Firewall Management Center makes a copy of the file in its repository. When you back up the database, these files are included. If you restore the database, the files are restored as well. To use a file in a file object, upload it through the interface, not directly to the file repository.

When you deploy configurations that specify a file object, the associated file is downloaded to the device in the appropriate directory.

You can take these actions for each file object:

- **Download**: Download the Secure Client file.
- **Edit**: Modify the file object details.
- **Delete**: Delete the Secure Client file object. When you delete a file object, the associated file is not deleted from the file repository, only the object is deleted.

Navigation path

Objects > **Object Management** > **VPN** > **Secure Client File**.

Fields

- **Name**: Enter the name of the file to identify the file object; you can add up to 128 characters.
- **File Name**: Click **Browse** to select the file. The file name and full path of the file are added when you select the file.
- **File Type**: Choose the file type corresponding to the file you have selected. These file types are available:

File Types	Details
Secure Client Image	Select when you add the Secure Client image you have downloaded from the Cisco Software Download Center. You can associate any new or additional Secure Client images to the remote access VPN policy. You can also remove unsupported or end-of-life client packages that are no longer required.
Secure Client VPN Profile	Choose for the Secure Client VPN profile file. The profile file is created using the GUI-based Secure Client Profile Editor, an independent configuration tool. See the *Secure Client Profile Editor* chapter in the appropriate release of the Cisco Secure Client (including AnyConnect) User Guide for details.

File Types	Details
Secure Client Management VPN Profile	Select when you add a profile file for the Secure Client management VPN tunnel. Download the Secure Client **VPN Management Tunnel Standalone Profile Editor** from Cisco Software Download Center if you have not done already and create a profile with required settings for the Secure Client management VPN tunnel.
AMP Enabler Service Profile	The profile is used for the Secure Client The AMP Enabler and this profile are deployed to endpoints from Firewall Threat Defense when a user connects to the VPN using remote access.
Feedback Profile	You can add a Customer Experience Feedback profile and select this type to receive information about the features and modules customers have enabled and use.
ISE Posture Profile	Choose if you are adding a profile file for the Secure Client ISE Posture module.
NAM Service Profile	Configure and add the NAM profile file using the Network Access Manager profile editor.
Network Visibility Service Profile	Profile file for Secure Client Network Visibility module. You can create the profile using the NVM profile editor.
Umbrella Roaming Security Profile	You must select this file type if you are deploying the Umbrella Roaming Security module using the .json file created using the profile editor.
Web Security Service Profile	Select this option when you add a profile file for the Web security module.
HostScan Posture Package	Select when you add a Secure Firewall Posture Package file. This file is used while configuring a Dynamic Access Policy (DAP) to collect information about the operating system, anti-virus, anti-spyware, and firewall software installed on the endpoints.
Secure Client External Browser Package	For selecting an external browser package file for SAML single sign-on web authentication. You can add the package file when a new version of the external package file is available. For more information, see Configure AAA Settings for Remote Access VPN, on page 1835.
Description	Add an optional description.

Related Topics

Cisco Secure Client Image, on page 1853
Group policy options for Secure Client, on page 3065



CHAPTER 104

Certificates

- Requirements and prerequisites for certificates, on page 3079
- Secure Firewall Threat Defense VPN certificate guidelines and limitations, on page 3079
- Manage Firewall Threat Defense certificates, on page 3080
- Install a certificate using self-signed enrollment, on page 3084
- Install a certificate using EST enrollment, on page 3084
- Install a certificate using SCEP enrollment, on page 3085
- Install a certificate using manual enrollment, on page 3086
- Install a certificate using a PKCS12 file, on page 3087
- Install a certificate using ACME enrollment, on page 3088
- Troubleshoot Threat Defense certificates, on page 3088

Requirements and prerequisites for certificates

This topic outlines the domain and role-based access requirements for certificate management.

Requirement	Supported values
Supported domains	Any
User roles	• Admin • Network Admin

Secure Firewall Threat Defense VPN certificate guidelines and limitations

Certificate enrollment guidelines

When a PKI enrollment object is associated with and then installed on a device, the certificate enrollment process starts immediately. The process is automatic for self-signed and SCEP enrollment types; it does not require any additional administrator's action. Manual certificate enrollment requires administrator's action.

- When a PKI enrollment object is associated with and then installed on a device, the certificate enrollment process starts immediately. The process is automatic for self-signed and SCEP enrollment types; it does not require any additional administrator's action. Manual certificate enrollment requires administrator's action.
- When importing identity certificates into Firewall Threat Defense devices, the expiration date of the certificate must not exceed February 6, 2106. Ensure that all certificates used within devices have expiration dates earlier to February 6, 2106 to avoid issues during certificate import or renewal.
- When the certificate enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your VPN Authentication Method.
- Firewall Threat Defense devices support certificate enrollment using Microsoft Certificate Authority(CA) Service, and CA Services provided on Cisco Adaptive Security Appliances(ASA) and Cisco IOS Router.
- Firewall Threat Defense devices cannot be configured as a certificate authority (CA).

Guidelines for certificates and multitenancy

Firewall Threat Defense devices support certificate enrollment using Microsoft CA Service, as well as CA services provided by Cisco Adaptive Security Appliances and Cisco IOS Routers.

Supported certificate enrollment services:

Certificate enrollment domains

- Certificate enrollment can be done in a **child domain** or a **parent domain**.
- When enrollment is done from a **parent domain**, the certificate enrollment object must also be in the same domain.
- If the trustpoint on a device is overridden in the child domain, the overridden value will be deployed on the device.

Enrollment visibility and behavior in leaf domains

- When certificate enrollment is done on a device in a **leaf domain**, the enrollment will be visible to the parent domain or another child domain.
- Adding additional certificates in a leaf domain is possible.

Certificate enrollment restrictions and permissions

- Once a device has certificates enrolled in one domain, it will be allowed to be enrolled in any other domain. The certificates can be added in the other domain.

Device movement between domains

- When you move a device from one domain to another, the certificates also get moved accordingly. You will receive an alert to delete the enrollments on these devices.

Manage Firewall Threat Defense certificates

This task enables you to manage certificates on devices including viewing certificate status, enrolling new certificates, and performing certificate operations such as export and re-enrollment.

See PKI infrastructure and digital certificates, on page 1718 for an introduction to Digital Certificates.

See Certificate enrollment objects, on page 2997 for a description of the objects used to enroll and obtain certificates on managed devices.

Procedure

Step 1 Select **Devices** > **Certificates**.

You can see the following columns for each device listed on this screen:

- **Name**—Lists the devices that already have trustpoints associated with them. Expand the device to see the list of associated trustpoints.
- **Domain**—Displays the certificates that are enrolled in a specific domain.
- **Enrollment Type**—Displays the type of enrollment used for a trustpoint.
- **Status**—Provides the status of the **CA Certificate** and **Identity Certificate**. You can view the certificate contents, when `Available`, by clicking the magnifying glass.

  When you view the CA certificate information, you can view the hierarchy of all the certifying authorities, which issued your CA certificate. If the enrollment fails, click status to view the failure message.
- Click **Enable weak-crypto** on the right to enable weak cipher usage in certificates. When you click the toggle button, you get a warning to confirm before enabling weak ciphers. Click **Yes** to enable weak ciphers.

  Note
  When a certificate enrollment fails due to weak cipher usage, you get a prompt to enable the weak cipher. You can choose to enable weak cipher when you need to use weak encryption.

- The additional column lists icons to perform the following tasks:
  - **Export Certificate**—Click to export and download a copy of the certificate. You can choose to export the PKCS12 (Complete Certificate Chain) or the PEM (Identity Certificate Only) format.

    You must provide a pass phrase to export a PKCS12 certificate format to import the file later.
  - **Re-enroll certificate**—Re-enroll an existing certificate.
  - **Refresh certificate status**—Refresh a certificate to synchronize the device certificate status to the Cloud-Delivered Firewall Management Center.
  - **Delete certificate**—Delete all the associated certificates for a trustpoint.

Step 2 Choose **(+) Add** to associate and install an enrollment object on a device.

When a certificate enrollment object is associated with and then installed on a device, the process of certificate enrollment starts immediately. The process is automatic for self-signed and SCEP enrollment types, meaning it does not require any additional administrator action. Manual certificate enrollment requires extra administrator action.

Note
The certificate enrollment on a device does not block the user interface and the enrollment process gets executed in the background, enabling the user to perform certificate enrollment on other devices in parallel.

The progress of these parallel operations can be monitored on the same user interface. The respective icons display the certificate enrollment status.

Related Topics

Install a certificate using self-signed enrollment, on page 3084
Install a certificate using SCEP enrollment, on page 3085
Install a certificate using manual enrollment, on page 3086
Install a certificate using a PKCS12 file, on page 3087

Automatically update CA bundles

You can set the management center to automatically update the CA certificates through CLI commands. By default, the CA certificates are automatically updated when you install or upgrade to version 7.0.5.



Note In an IPv6-only deployment, the automatic update of CA certificates may fail, because, some of the Cisco servers do not support IPv6. In such cases, force update the CA certificates using the **configure cert-update run-now force** command.

Procedure

Step 1 Log into the FMC CLI using SSH, or, if virtual, open the VM console.

Step 2 You can verify whether the CA certificates in the local system are the latest or not:

configure cert-update test

This command compares the CA bundle on the local system with the latest CA bundle (from the Cisco server). If the CA bundle is up to date, no connection check is executed and the test result is displayed as the one below:

Example:

```
> configure cert-update test
Test succeeded, certs can safely be updated or are already up to date.
```

If the CA bundle is out of date, the connection check is executed on the downloaded CA bundle and the test result is displayed.

Example:

When the connection check fails:

```
> configure cert-update test
Test failed, not able to fully connect.
```

Example:

When the connection check succeeds, or the CA bundle is already up to date:

```
> configure cert-update test
Test succeeded, certs can safely be updated or are already up to date.
```

Step 3 (Optional) To instantly update the CA bundles:

configure cert-update run-now

Example:

```
>configure cert-update run-now
Certs have been replaced or was already up to date.
```

When you execute this command, the CA certificates (from the Cisco server) are verified for SSL connectivity. If the SSL connectivity check fails for even one of the Cisco servers, the process is terminated.

Example:

```
> configure cert-update run-now
Certs failed some connection checks.
```

To proceed with the update despite connection failures, use the **force** keyword.

Example:

```
> configure cert-update run-now force
Certs failed some connection checks, but replace has been forced.
```

Step 4 If you do not want the CA bundles to be automatically updated, disable the configuration:

configure cert-update auto-update disable

Example:

```
> configure cert-update auto-update disable
Autoupdate is disabled
```

Step 5 To re-enable the automatic update of CA bundles:

configure cert-update auto-update enable

Example:

```
> configure cert-update auto-update enable
Autoupdate is enabled and set for every day at 12:18 UTC
```

When you enable the automatic update on the CA certificates, the update process is executed daily at a system-defined time.

Step 6 (Optional) View the status of automatic update of CA certificates:

show cert-update

Example:

```
> show cert-update
```

```
Autoupdate is enabled and set for every day at 09:34 UTC
CA bundle was last modified 'Thu Sep 15 16:12:35 2022'
```

Install a certificate using self-signed enrollment

Install a self-signed certificate to establish secure VPN connections by creating an identity certificate directly on the managed device.

Procedure

Step 1 On the **Devices** > **Certificates** screen, choose **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop-down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the type Self-Signed from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Add certificate enrollment objects, on page 2999.

Step 4 Press **Add** to start the Self Signed, automatic, enrollment process.

For self signed enrollment type trustpoints, the **CA Certificate** status will always be displayed, since the managed device is acting as its own CA and does not need a CA certificate to generate its own Identity Certificate.

The **Identity Certificate** will go from InProgress to Available as the device creates its own self signed identity certificate.

Step 5 Click the magnifying glass to view the self-signed Identity Certificate created for this device.

What to do next

When enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your Site to Site and Remote Access VPN Authentication Method.

Install a certificate using EST enrollment

Use EST enrollment to automatically obtain and install identity certificates on devices from a Certificate Authority (CA).

Before you begin

Note Using EST enrollment establishes a direct connection between the managed device and the CA server. So be sure your device is connected to the CA server before beginning the enrollment process.

Note EST's ability to auto-enroll a device when its certificate expires is not supported.

Follow these steps to install a certificate using EST enrollment:

Procedure

Step 1 On the **Devices** > **Certificates** screen, click **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop-down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose the EST certificate enrollment object from the **Cert Enrollment** drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Add certificate enrollment objects, on page 2999.

Step 4 Click **Add** to enroll the certificate on the device.

The **Identity Certificate** will go from `InProgress` to `Available` as the device obtains its identity certificate using EST from the specified CA. Sometimes, a manual refresh might be required to obtain the identity certificate.

Step 5 Click the magnifying glass to view the Identity Certificate created and installed on this device.

Install a certificate using SCEP enrollment

Install a certificate using SCEP enrollment to enable secure authentication for VPN connections and other security features on your device.

Before you begin

Note Using SCEP enrollment establishes a direct connection between the managed device and the CA server. So be sure your device is connected to the CA server before beginning the enrollment process.

Follow these steps to install a certificate using SCEP enrollment:

Procedure

Step 1 On the **Devices** > **Certificates** screen, choose **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop-down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the type SCEP from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Add certificate enrollment objects, on page 2999.

Step 4 Press **Add**, to start the automatic enrollment process.

For SCEP enrollment type trustpoints, the **CA Certificate** status will transition from InProgress to Available as the CA Certificate is obtained from the CA server and installed on the device.

The **Identity Certificate** will go from `InProgress` to `Available` as the device obtains its identity certificate using SCEP from the specified CA. Sometimes, a manual refresh might be required to obtain the identity certificate.

Step 5 Click the magnifying glass to view the Identity Certificate created and installed on this device.

What to do next

When enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your Site to Site and Remote Access VPN Authentication Method.

Install a certificate using manual enrollment

Manual enrollment allows you to obtain and install an identity certificate on a managed device by generating a certificate signing request (CSR) and working with your PKI CA server to obtain the corresponding identity certificate.

Procedure

Step 1 On the **Devices** > **Certificates** screen, choose **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop-down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the type Manual from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Add certificate enrollment objects, on page 2999.

Step 4 Press **Add** to start the enrollment process.

Step 5 Execute the appropriate activity with your PKI CA Server to obtain an identity certificate.

a) Click **Identity Certificate** warning to view and copy the CSR.
b) Execute the appropriate activity with your PKI CA Server to obtain an identity certificate using this CSR.

   This activity is completely independent of the Cloud-Delivered Firewall Management Center or the managed device. When complete, you will have an Identity Certificate for the managed device. You can place it in a file.

c) To finish the manual process, install the obtained identity certificate onto the managed device.

   Return to the Cloud-Delivered Firewall Management Center dialog and select **Browse Identity Certificate** to choose the identity certificate file.

   Note
   Ensure not to choose a binary certificate (PKCS12, DER, and alike) file because Firewall Threat Defense does not support them.

Step 6 Select **Import** to import the Identity Certificate.

The Identity Certificate status will be `Available` when the import complete.

Step 7 Click the magnifying glass to view the **Identity Certificate** for this device.

What to do next

When enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your Site to Site and Remote Access VPN Authentication Method.

Install a certificate using a PKCS12 file

Install a certificate using a PKCS12 file to enable secure authentication and encrypted communications on managed devices.

Procedure

Step 1 Go to **Devices** > **Certificates** screen, choose **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a pre-configured managed device from the **Device** drop down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the PKCS type from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Add certificate enrollment objects, on page 2999.

Step 4 Press **Add**.

The CA Certificate and Identity Certificate status will go from `In Progress` to `Available` as it installs the PKCS12 file on the device.

Note

When you upload the PKCS12 file for the first time, the file is stored in Cloud-Delivered Firewall Management Center as part of the CertEnrollment object. For any failed enrollments due to a wrong passphrase or failed deployment, retry enrolling the PKCS12 certificate without uploading the file again. Also, whenever you modify the certificate enrollment object to permit overrides, you must update the **Passphrase** for the certificate for a sucessful enrollment. A PKCS12 file size should not be larger than 24K.

Step 5 Once `Available`, click the magnifying glass to view the Identity Certificate for this device.

What to do next

The certificate (trustpoint) on the managed device is named the same as the PKCS#12 file. Use this certificate in your VPN authentication configuration.

Install a certificate using ACME enrollment

Install certificates automatically using the ACME (Automatic Certificate Management Environment) protocol to eliminate manual certificate installation and renewal processes.

Procedure

Step 1 Choose **Devices** > **Certificates**.

Step 2 Click **Add** to open the **Add New Certificate** dialog.

Step 3 Choose a device from the **Device** drop-down list.

Step 4 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the type ACME from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Add certificate enrollment objects, on page 2999.

Step 5 Click **Add**, to start the automatic enrollment process.

For ACME enrollment type trustpoints, the status shows `InProgress` until the enrollment is successful. CA is not applicable for ACME.

Step 6 Click the magnifying glass to view the Identity Certificate created and installed on this device.

What to do next

When enrollment is complete, a trustpoint is configured on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your remote access VPN authentication.

Troubleshoot Threat Defense certificates

See Secure Firewall Threat Defense VPN certificate guidelines and limitations, on page 3079 to determine if variations in your certificate enrollment environment may be causing a problem.

Certificate expiration monitoring CA server connectivity

Ensure there is a route to the CA Server from the device. If the CA Server's host name is given in the Enrollment Object, use Flex Config to configure DNS appropriately to reach the server. Alternatively, use the IP Address of the CA Server.

Microsoft 2012 CA server template configuration

If you are using a Microsoft 2012 CA Server, the default IPsec Template is not accepted by the managed device and must be changed.

To configure a working template, follow these steps as you use MS CA documentation as a reference.

1. Duplicate the IPsec (Offline Request) template.
2. In **Extensions** > **Application policies**, select *IP security end system*, instead of the *IP security IKE intermediate*.
3. Set the permissions and the template name.
4. Add the new template and change the registry settings to reflect the new template name.

Default keyring certificate regeneration

On the Cloud-Delivered Firewall Management Center, you might receive this health alert related to the Firewall Threat Defense device:

```
Code - F0853; Description - default Keyring's certificate is invalid, reason: expired
```

Solution: In such cases, use the following command to regenerate the default certificate in CLISH CLI:

```
> system support regenerate-security-keyring default
```

CA certificate configuration errors

A red cross appears in the CA certificate status with this error:

Error: `Fail to configure CA certificate`

Solution: See Troubleshoot Certificate Error on FMC.

Certificate verification in .pfx files

To check the list of the certificates in a .pfx file, use tools such as `certutil` or `openssl`. You can see the whole chain with ID certificate, SubCA certificate, and CA certificate (if any).

- `certutil -dump cert.pfx`
- `openssl pkcs12 -info -in cert.pfx`

Identity certificate import errors

Error: `Identity certificate import required`

Solution: See Troubleshoot Certificate Error "Identity certificate import required" on FMC.

Self-signed certificates hostname behavior in Threat Defense high-availability environments

In a Threat Defense HA configuration, the secondary unit adopts the primary unit's hostname. When you break the HA pair, you must perform these actions:

- Manually update the hostname on the secondary unit to its unique identifier
- Perform a certificate re-enrollment to ensure that the certificate accurately reflects the unit's identity.



PART XXIII

Reference

- FAQ and Support, on page 3093
- Cloud-Delivered Firewall Management Center command line reference, on page 3109
- Communication Ports and Internet Access, on page 3115



CHAPTER 105

FAQ and Support

- **Security Cloud Control Firewall Management Platform Maintenance Schedule**, on page 3093
- Navigate from Security Cloud Control to Cloud-Delivered Firewall Management Center, on page 3094
- What does the default action "Analyze all tunnel traffic" for prefiltering mean?, on page 3094
- How Security Cloud Control processes personal information, on page 3094
- Can I restore a backup from a different device?, on page 3094
- Does deploying a new prefilter policy immediately affect ongoing sessions?, on page 3095
- How do I keep my security databases and feeds up to date?, on page 3095
- What version of Secure Firewall Threat Defense can I manage with Cloud-Delivered Firewall Management Center?, on page 3095
- How do I exclude specific traffic (Webex, Zoom, etc) from the remote access VPN?, on page 3096
- How do I prevent users from accessing undesirable external network resources, such as inappropriate websites?, on page 3097
- Security Feed Questions, on page 3097
- How do I setup Rate-Based Attack Prevention on the FTD using Snort 2?, on page 3100
- Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI, on page 3101
- About Data Interfaces, on page 3105
- End-of-Support for management of Secure Firewall Threat Defense devices, Version 7.0.*x*, managed by Cloud-Delivered Firewall Management Center, on page 3106

Security Cloud Control Firewall Management Platform Maintenance Schedule

Security Cloud Control Firewall Management updates its platform every week with new features and quality improvements. Updates are made during a 3-hour period according to this schedule:

Day of the week	Time of day (24-hour time, UTC)
Thursday	09:00 UTC - 12:00 UTC

During the Security Cloud Control Firewall Management upgrade period, you can still access your organization and the Cloud-Delivered Firewall Management Center. Additionally, the devices that you have onboarded to Security Cloud Control Firewall Management continue to enforce their security policies.

Note

- During the maintenance period of the Security Cloud Control Firewall Management, refrain from creating or deploying configurations to the devices it manages.
- If there is any issue that stops Security Cloud Control Firewall Management from communicating, Cisco addresses that issue in all the affected tenants as quickly as possible, even if it is outside the maintenance window.

Navigate from Security Cloud Control to Cloud-Delivered Firewall Management Center

1. In the Security Cloud Control dashboard, click **Administration** > **Integrations** > **Firewall Management Center**.
2. With Cloud-Delivered Firewall Management Center selected, click **Devices** in the Management pane.

What does the default action "Analyze all tunnel traffic" for prefiltering mean?

“Analyze all tunnel traffic” means subject all network traffic to the rules in the access control policy after they have been analyzed by the prefilter policy.

How Security Cloud Control processes personal information

To learn how Security Cloud Control processes your personal identifiable information, see the Security Cloud Control Privacy Data Sheet.

Can I restore a backup from a different device?

Yes, if the devices are the same model, are running the same software version, have the same number of network modules, and the same number of physical interfaces.

Does deploying a new prefilter policy immediately affect ongoing sessions?

No. When you deploy a prefilter policy, its rules are not applied on the existing tunnel sessions. Hence, traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic on an existing connection that could match a policy is omitted from the hit count.

How do I keep my security databases and feeds up to date?

If the management center has internet access, the system can often obtain updates on security databases and feeds directly from Cisco. We recommend you schedule or enable automatic content updates whenever possible. Some updates are auto-enabled by the initial setup process or when you enable the related feature. Other updates you must schedule yourself. After initial setup, we recommend you review all auto-updates and adjust them if necessary.

- There are several security databases and feeds you should update:
- Vulnerability database (VDB)
- Geolocation database (GeoDB)
- Intrusion rules (SRU/LSP)
- Security Intelligence Feeds
- URL categories and reputations

What version of Secure Firewall Threat Defense can I manage with Cloud-Delivered Firewall Management Center?

Cloud-Delivered Firewall Management Center supports these versions of Secure Firewall Threat Defense:

- Verion 7.0.3 or later 7.0.x versions.
- Version 7.2 and later versions.

Note Software Version 7.1 is _not_ supported.

All hardware and virtual deployments which can run these software versions are supported.

How do I exclude specific traffic (Webex, Zoom, etc) from the remote access VPN?

You can exclude specific traffic from the remote access VPN using dynamic split tunneling based on DNS domain names.

Excluded domains are not blocked. Instead, traffic to those domains is kept outside the VPN tunnel. For example, you could send traffic to Cisco WebEx on the public Internet, thus freeing bandwidth in your VPN tunnel for traffic that is targeted to servers within your protected network.

Procedure

Step 1 From the Security Cloud Control home page, in the navigation bar, click **Security Devices**.

Step 2 Find the Secure Firewall Threat Defense device you want to add this rule to. You can use the filter or search field to find the device.

Step 3 Select the device, and in the Device Management pane at the right, click **Device Overview**.

Step 4 Configure the group policy to use Dynamic Split Tunnel.

a) Choose **Devices** > **Remote Access**.
b) Click **Edit** on the remote access VPN policy for which you want to configure dynamic split tunneling.
c) Click **Edit** on the required connection profile.
d) Click **Edit Group Policy**.

Step 5 Configure the Secure Client custom attribute in the Add/Edit Group Policy dialog box.

a) Click the Secure Client tab.
b) Click **Custom Attributes** and click +.
c) Choose **Dynamic Split Tunneling** from the Secure Client Attribute drop-down list.
d) Click + to create a new custom attribute object.
e) Enter the name for the custom attribute object.
f) Exclude domains—Specify domain names that will be excluded from the remote access VPN.
g) Click **Save**.
h) Click **Add**.

Step 6 Verify the configured custom attribute and click **Save**.

Step 7 When you are ready to deploy this change to the device, click **Deploy** in the menu bar at the top of the page.

How do I prevent users from accessing undesirable external network resources, such as inappropriate websites?

Procedure

Step 1 From the Security Cloud Control home page, click **Security Devices** in the navigation bar.

Step 2 Find the Secure Firewall Threat Defense device you want to add these rules to. You can use the filter or search field to find the device.

Step 3 Select the device and in the Policies pane at the right, click **Access Control**.

Step 4 Click the policy you want to update.

Step 5 Click **Add Rule**.

Step 6 Give the rule a name.

Step 7 In the **Action** field, select **Block**.

Step 8 Insert the rule into either the Mandatory or into the Default policy.

Step 9 Click the **URLs** tab.

Step 10 In the Categories section, check the categories you want to block and accept the default value for "Any Reputation".

Step 11 Click **Add URL**.

Step 12 If there are specific URLs that you want to block, you can do that by entering them in the **Manually Enter URL** field and then click **Add URL**.

Step 13 Click **Apply**.

Step 14 On the policy page, click **Save**.

Step 15 When you are ready to deploy this change to the device, click **Deploy** in the menu bar at the top of the page.

Note
Note: This instruction assumes you have the URL filtering license

Security Feed Questions

Related Information

How do I update intrusion rules (SRU/LSP)?

Follow this procedure to configure recurring Intrusion Rule update downloads.

Procedure

Step 1 From the Cloud-Delivered Firewall Management Center home page, navigate **System (gear icon) > Updates > Rule Updates**.

Step 2 Under **Recurring Rule Update Imports**, check **Enable Recurring Rule Update Imports**.

Step 3 Specify the **Import Frequency** and start time.

Note
As updates are published multiple times each week, it is recommended to check on a Daily basis.

Step 4 (Optional, but recommended) Check **Reapply all policies...** to deploy after each update.

Caution
Deploying intrusion rule updates can cause a Snort restart in rare occasions. It is recommended to deploy intrusion rule updates during a maintenance window.

Step 5 Click **Save**.

Step 6 Deploy your changes when you are ready.

Caution
Deploying intrusion rule updates can cause a Snort restart. We recommend you deploy intrusion rule updates during a maintenance window.

How do I update my Cisco vulnerability database (VDB)?

The initial setup on the management center automatically downloads and installs the latest VDB from Cisco as a one-time operation. It also schedules a weekly task to download the latest available software updates, which includes the latest VDB. We recommend you review this weekly task and adjust if necessary, by navigating in the Cloud-Delivered Firewall Management Center to **System gear icon** > **Tools** > **Scheduling**.

You can choose to update the VDB either manually or automate the download and updates by creating scheduled tasks in the system. We recommend that you choose the automated method, to ensure timely downloads and updates.

Before you begin

.

You must be in the global domain to perform this task.

Procedure

Step 1 To manually update the VDB, see Manually Update the VDB and follow the steps in the procedure.

Step 2 To automate VDB updates, see Vulnerability Database Update Automation. Ensure you create tasks for both download of VDB versions and installation.

How do I update my Geolocation database?

As a part of initial configuration, the system configures a weekly automatic Geolocation (GeoDB) update. If configuring the update fails, we recommend you configure regular GeoDB updates as described in this procedure.

Procedure

Step 1 From the Cloud-Delivered Firewall Management Center home page, **System (gear icon) > Updates > Updates > Geolocation Updates**.

Step 2 Under **Recurring Geolocation Updates**, check **Enable Recurring Weekly Updates from the Support Site**.

Step 3 Specify the **Update Start Time**.

Step 4 Click **Save**.

How do I update Security Intelligence feeds?

By default, the built-in feeds on Cloud-Delivered Firewall Management Center are updated every two hours, with updates being pushed immediately to managed devices.

To change the update configuration, perform the following steps:

Procedure

Step 1 From the Cloud-Delivered Firewall Management Center home page, navigate **Objects > Object Management**.

Step 2 Expand the Security Intelligence node, then choose the feed type whose frequency you want to change.

Step 3 Next to the feed you want to update, click the pencil icon to **Edit** the update frequency.

Note
The system-provided URL feed is combined with the domain feed under DNS Lists and Feeds.

Step 4 Edit the **Update Frequency**.

Step 5 Click **Save**.

How do I update URL reputations?

If you Enable Automatic Updates, by default, automatic URL updates are enabled. The management center checks Talos updates every 30 minutes. If you need strict control over when the system contacts external

resources, disable automatic updates and instead create a recurring task using the scheduler. Although daily updates tend to be small, if it has been more than five days since your last update, new URL filtering data may take up to 20 minutes to download, depending on your bandwidth. Then, it may take up to 30 minutes to perform the update itself.

Procedure

Step 1 From the Cloud-Delivered Firewall Management Center home page, navigate **Integration > Other Integrations**.

Step 2 Click **Cloud Services**.

Step 3 In the URL Filtering pane:

a) Enable URL filtering.

b) Enable automatic updates.

Step 4 Click **Save**.

How do I setup Rate-Based Attack Prevention on the FTD using Snort 2?

Dynamic rule states are policy-specific.

A **Revert** appears in a field when you enter an invalid value; click it to revert to the last valid value for that field or to clear the field if there was no previous value.

Note Dynamic rule states cannot enable disabled rules or drop traffic that matches disabled rules.

Procedure:

Procedure

Step 1 On the Security Cloud Control menu bar, click **Administration** > **Integrations** > **Firewall Management Center** to view the Services page.

Step 2 Choose Cloud-Delivered FMC and click the links in the Actions, Management, or System pane to navigate to Cloud-Delivered Firewall Management Center to perform various actions. See View Services Page Information.

Step 3 Choose **Policies > Access Control > Intrusion**.

Step 4 Click Snort 2 Version next to the policy you want to edit.

If View (View button) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 5 Click **Rules** immediately under **Policy Information** in the navigation pane.

Step 6 Choose the rule or rules where you want to add a dynamic rule state.

Step 7 Choose **Dynamic State > Add Rate-Based Rule State**.

Step 8 Choose a value from the Track By drop-down list.

Step 9 If you set Track By to Source or Destination, enter the address of each host you want to track in the Network field. You can specify a single IP address, address block, variable, or a comma-separated list comprised of any combination of these.

Step 10 Next to Rate, specify the number of rule matches per time period to set the attack rate:

Step 11 From the New State drop-down list, specify the new action to be taken when the conditions are met.

Step 12 Enter a value in the Timeout field.

After the timeout occurs, the rule reverts to its original state. Specify 0 or leave the Timeout field blank to prevent the new action from timing out.

Step 13 Click OK.

Note
The system displays a Dynamic State next to the rule in the Dynamic State column. If you add multiple dynamic rule state filters to a rule, a number over the filter indicates the number of filters.

Note
To delete all dynamic rule settings for a set of rules, choose the rules on the Rules page, then choose Dynamic State > Remove Rate-Based States. You can also delete individual rate-based rule state filters from the rule details for the rule by choosing the rule, clicking Show details, then clicking Delete by the rate-based filter you want to remove.

Step 14 To save changes you made in this policy since the last policy commit, click Policy Information, then click Commit Changes.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

Complete the initial configuration of a Secure Firewall Threat Defense device using the CLI

Connect to the device's CLI to perform initial setup, including setting the management IP address, gateway, and other basic networking settings using the setup wizard. Ensure all DNS and firewall ports are accessible for communication.

The dedicated management interface is a special interface with its own network settings. If you do not want to use the management interface, you can use the CLI to configure a data interface instead.

This configuration is ideal for devices that are going to be onboarded with their CLI registration key.

Note Do **not** use this configuration procedure for devices that are onboarding with zero-touch provisioning.

Procedure

Step 1 Connect to the device's CLI, either from the console port or using SSH to the management interface. If you intend to change the network settings, we recommend using the console port so you do not get disconnected.

(Firepower and Secure Firewall hardware models) The console port connects to the FXOS CLI. The SSH session connects directly to the Firewall Threat Defense CLI.

Step 2 Log in with the username **admin** and the password **Admin123** .

At the console port, you connect to the FXOS CLI. The first time you log in to FXOS, you are prompted to change the password. This password is also used for the Firewall Threat Defense login for SSH.

Note
If the password was already changed, and you do not know it, you must reimage the device to reset the password to the default.

For Firepower and Secure Firewall hardware, see the Reimage Procedures in the Cisco FXOS Troubleshooting Guide for the Firewall Threat Defense .

For the ISA 3000, see the Cisco Secure Firewall ASA and Threat Defense Reimage Guide .

Example:

```
firepower login:
                          admin
Password:
                          Admin123
Successful login attempts for user 'admin' : 1
[...]
Hello admin. You must change your password.
Enter new password:
                          ********
Confirm new password:
                          ********
Your password was updated successfully.
[...]
firepower#
```

Step 3 If you connected to FXOS on the console port, connect to the Firewall Threat Defense CLI.

connect ftd

Example:

```
firepower# connect ftd
>
```

Step 4 The first time you log in to the device, you are prompted to accept the End User License Agreement (EULA) and, if using an SSH connection, to change the admin password. You are then presented with the CLI setup script.

Note
You cannot repeat the CLI setup wizard unless you clear the configuration; for example, by reimaging. However, all of these settings can be changed later at the CLI using **configure network** commands. See the threat defense command reference.

Defaults or previously entered values appear in brackets. To accept previously entered values, press **Enter**.

Note
The management interface settings are used even when you enable Firewall Threat Defense access on a data interface. For example, the management traffic that is routed over the backplane through the data interface will resolve FQDNs using the management interface DNS servers, and not the data interface DNS servers.

See the following guidelines:

- **Configure IPv4 via DHCP or manually?**—If you want to use a data interface for Firewall Threat Defense access instead of the management interface, choose **manual**. Although you do not plan to use the management interface, you must set an IP address, for example, a private address. You cannot configure a data interface for management if the management interface is set to DHCP, because the default route, which must be **data-interfaces** (see the next bullet), might be overwritten with one received from the DHCP server.
- **Enter the IPv4 default gateway for the management interface**—If you want to use a data interface for Firewall Threat Defense access instead of the management interface, set the gateway to be **data-interfaces**. This setting forwards management traffic over the backplane so it can be routed through the FMC access data interface.
- **If your networking information has changed, you will need to reconnect**—If you are connected with SSH but you change the IP address at initial setup, you will be disconnected. Reconnect with the new IP address and password. Console connections are not affected.
- **Manage the device locally?**—Enter **YES** to configure the device for the device to be managed by either the Cloud-Delivered Firewall Management Center or Secure Firewall Device Manager.

  Manage the device locally?—Enter **NO** to configure the device for remote management with the on-premises Firewall Management Center.
- **Configure firewall mode?**—We recommend that you set the firewall mode at initial configuration. Changing the firewall mode after initial setup erases your running configuration. Note that data interface Firewall Threat Defense access is only supported in routed firewall mode.

Step 5 (Optional) Configure a data interface for Cloud-Delivered Firewall Management Center access.

configure network management-data-interface

You are then prompted to configure basic network settings for the data interface.

Note
You should use the console port when using this command. If you use SSH to the Management interface, you might get disconnected and have to reconnect to the console port. See below for more information about SSH usage.

See the following details for using this command. See About Data Interfaces, on page 3105 for more informatio.

- The original management interface cannot use DHCP if you want to use a data interface for management. If you did not set the IP address manually during initial setup, you can set it now using the **configure network** {**ipv4** | **ipv6**} **manual** command. If you did not already set the management interface gateway to **data-interfaces**, this command will set it now.
- When you onboard the device for Firewall Threat Defense management through Security Cloud Control, Security Cloud Control discovers and maintains the interface configuration, including the following settings: interface name and IP address, static route to the gateway, DNS servers, and DDNS server. For more information about the DNS server configuration, see below. You can later make changes to the

access interface configuration, but make sure you don't make changes that can prevent the device or Security Cloud Control from re-establishing the management connection. If the management connection is disrupted, the device includes the **configure policy rollback** command to restore the previous deployment.

- This command sets the *data* interface DNS server. The Management DNS server that you set with the setup script (or using the **configure network dns servers** command) is used for management traffic. The data DNS server is used for DDNS (if configured) or for security policies applied to this interface.

  Also, local DNS servers are only retained if the DNS servers were discovered at initial registration. For example, if you registered the device using the Management interface, but then later configure a data interface using the **configure network management-data-interface** command, then you must manually configure all of these settings in Security Cloud Control, including the DNS servers, to match the device configuration.
- You can change the management interface after you onboard the Firewall Threat Defense for Firewall Threat Defense management through threat defense, to either the management interface or another data interface.
- The FQDN that you set in the setup wizard will be used for this interface.
- You can clear the entire device configuration as part of the command; you might use this option in a recovery scenario, but we do not suggest you use it for initial setup or normal operation.
- To disable data managemement, enter the **configure network management-data-interface disable** command.

Example:

```
> configure network management-data-interface
Data interface to use for management: ethernet1/1
Specify a name for the interface [outside]:
IP address (manual / dhcp) [dhcp]:
DDNS server update URL [none]:
https://jcrichton:pa$$w0rd17@domains.example.com/nic/update?hostname=<h>&myip=<a>
Do you wish to clear all the device configuration before applying ? (y/n) [n]:

Configuration done with option to allow FMC access from any network, if you wish to change
 the FMC access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.

>
```

Example:

```
> configure network management-data-interface
Data interface to use for management: ethernet1/1
Specify a name for the interface [outside]: internet
IP address (manual / dhcp) [dhcp]: manual
IPv4/IPv6 address: 10.10.6.7
Netmask/IPv6 Prefix: 255.255.255.0
Default Gateway: 10.10.6.1
Comma-separated list of DNS servers [none]: 208.67.222.222,208.67.220.220
DDNS server update URL [none]:
Do you wish to clear all the device configuration before applying ? (y/n) [n]:
```

```
Configuration done with option to allow FMC access from any network, if you wish to change
 the FMC access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.

>
```

Step 6 (Optional) Limit data interface access to Security Cloud Control on a specific network.

configure network management-data-interface client *ip_address netmask*

By default, all networks are allowed.

About Data Interfaces

You can use either the dedicated management interface or a regular data interface for communication with the device. Security Cloud Control access on a data interface is useful if you want to manage the FTD remotely from the outside interface, or you do not have a separate management network. Security Cloud Control supports high availability on the FTD managed remotely from the data interface.

FTD management access from a data interface has the following limitations:

- You can only enable manager access on one physical, data interface. You cannot use a subinterface or EtherChannel.
- Routed firewall mode only, using a routed interface.
- PPPoE is not supported. If your ISP requires PPPoE, you will have to put a router with PPPoE support between the FTD and the WAN modem.
- The interface must be in the global VRF only.
- SSH is not enabled by default for data interfaces, so you will have to enable SSH later using Security Cloud Control. Because the management interface gateway will be changed to be the data interfaces, you also cannot SSH to the management interface from a remote network unless you add a static route for the management interface using the **configure network static-routes** command. For FTDv on Amazon Web Services, a console port is not available, so you should maintain your SSH access to the management interface: add a static route for Management before you continue with your configuration. Alternatively, be sure to finish all CLI configuration (including the **configure manager add** command) before you configure the data interface.

End-of-Support for management of Secure Firewall Threat Defense devices, Version 7.0.*x*, managed by Cloud-Delivered Firewall Management Center

Which version of Secure Firewall Threat Defense devices has Cloud-Delivered Firewall Management Center stopped managing as of December 15, 2025?

Cloud-Delivered Firewall Management Center stopped managing Secure Firewall Threat Defense devices, Version 7.0.*x*, as of December 15, 2025.

Why is this necessary?

This was necessary to ensure compatibility and security for our customers' systems. As we upgrade our systems, it was essential that only the latest version of Secure Firewall Threat Defense devices were supported to maintain reliability and performance.

Can I onboard Version 7.0.*x* devices after December 15, 2025?

No, you can no longer onboard Secure Firewall Threat Defense, Version 7.0.*x*, devices after December 15, 2025.

After December 15, 2025, will I be able to use my existing Secure Firewall Threat Defense 7.0.*x* devices?

Cloud-Delivered Firewall Management Center will not support the management of Secure Firewall Threat Defense devices running Version 7.0.*x* after December 15, 2025. You will not be able to make or deploy changes to Secure Firewall Threat Defense devices running Version 7.0.*x* and earlier except to upgrade or unregister.



Note Secure Firewall Threat Defense devices running Version 7.0.*x* will still be visible in Cloud-Delivered Firewall Management Center, but in read-only mode only. Changes you make after December 15, 2025 will not be deployed to the devices.

What are the recommended upgrade options?

You are required to upgrade to at least Version 7.2.*x*. The Cisco-suggested version is 7.6.2.

How do I upgrade?

For information about upgrading Secure Firewall Threat Defense devices, refer to Threat defense upgrade guide for Cloud-Delivered Firewall Management Center.

Are there costs associated with upgrade?

There is no addtional cost, unless your hardware must also be upgraded. (The reference to hardware is limited to Cisco Secure ASA 5508-X and Cisco Secure ASA 5516-X running Secure Firewall Threat Defense Version 7.0.*x*). If you must refresh your hardware, reach out to your Cisco account team.

What will happen if I don't upgrade?

From December 15, 2025, you will not be able to make or deploy changes to Secure Firewall Threat Defense devices running Version 7.0.*x* and earlier, except to upgrade or unregister. Failing to upgrade leaves you vulnerable to compatibility issues, security issues, and possible downtime. To minimize disruption, you are required to upgrade now.

Will there be any impact to my existing systems or applications?

Upgrading is unlikely to have any significant impact on your current deployment. If you face any issues during the upgrade process, contact the Cisco Technical Assistance Center (TAC).

For assistance in upgrading, Cisco offers the LevelUp program. You can request assistance using this link.

Can I migrate my existing ASA 5508 and 5516 devices to Cloud-Delivered Firewall Management Center?

We recommend that you do *not* migrate your existing ASA 5508 and ASA 5516 devices to Cloud-Delivered Firewall Management Center because the last supported Secure Firewall Threat Defense version for these ASA devices is 7.0.

In addition, Cisco has stopped support for management of Secure Firewall Threat Defense Version 7.0.x in Cloud-Delivered Firewall Management Center as of December 31, 2025, and the ASA 5508 and 5516 devices cannot be upgraded to Version 7.2.x and later.

Can I manage my device's health in read-only mode after support ends?

After December 15, 2025, you will not be able to deploy device health policy changes, if any, to Secure Firewall Threat Defense devices running Version 7.0.x and earlier. These devices will continue to operate with the previously configured health policies.

When will Cloud-Delivered Firewall Management Center end support for management of Version 7.2.*x*?

Cloud-Delivered Firewall Management Center will support the management of Secure Firewall Threat Defense devices running Version 7.2.*x* until October 31, 2026.

What if I'm unsure about how to proceed?

Reach out to Cisco Technical Assistance Center (TAC) (or the Cisco account team) for clarifications.



CHAPTER 106

Cloud-Delivered Firewall Management Center command line reference

This reference explains the command line interface (CLI) for the Cloud-Delivered Firewall Management Center.

Note For Secure Firewall Threat Defense, see the *Cisco Secure Firewall Threat Defense Command Reference*.

- About the Cloud-Delivered Firewall Management Center CLI, on page 3109
- Cloud-Delivered Firewall Management Center CLI management commands , on page 3110
- Cloud-Delivered Firewall Management Center CLI show commands, on page 3111
- Cloud-Delivered Firewall Management Center CLI configuration commands, on page 3112
- Cloud-Delivered Firewall Management Center CLI system commands, on page 3112

About the Cloud-Delivered Firewall Management Center CLI

The Cloud-Delivered Firewall Management Center CLI provides access to view and troubleshoot your Cloud-Delivered Firewall Management Center.

When you use SSH to log into the Cloud-Delivered Firewall Management Center, you access the CLI. You can then access the Linux shell using the **expert** command, but this is strongly discouraged.

Caution We strongly recommend that you do not access the Linux shell unless directed by Cisco TAC or explicit instructions in the Secure Firewall user documentation.

Caution Users with Linux shell access can obtain root privileges, which can present a security risk. For system security reasons, we strongly recommend:

- If you establish external authentication, make sure that you restrict the list of users with Linux shell access appropriately.
- Do not establish Linux shell users in addition to the pre-defined **admin** user.

You can use the commands described in this appendix to view and troubleshoot your Cloud-Delivered Firewall Management Center. You can also perform limited configuration operations.

Cloud-Delivered Firewall Management Center CLI modes

The CLI encompasses four modes. The default mode, CLI Management, includes commands for navigating within the CLI itself. The remaining modes contain commands addressing three different areas of Cloud-Delivered Firewall Management Center functionality; the commands within these modes begin with the mode name: `system`, `show`, or `configure`.

When you enter a mode, the CLI prompt changes to reflect the current mode. For example, to display version information about system components, you can enter the full command at the standard CLI prompt:

```
> show version
```

If you have previously entered **show** mode, you can enter the command without the **show** keyword at the show mode CLI prompt:

```
show> version
```

Cloud-Delivered Firewall Management Center CLI management commands

The CLI management commands provide the ability to interact with the CLI. These commands do not affect the operation of the device.

exit command

This command moves the CLI context up to the next highest CLI context level. Issuing this command from the default mode logs the user out of the current CLI session.

Syntax: `exit`

Example:

```
system> exit
>
```

expert command

This command invokes the Linux shell.

Syntax: `expert`

Example:

```
> expert
```

? (question mark) command

This command displays context-sensitive help for CLI commands and parameters. Use the question mark (`?`) command in these ways:

- To display help for the commands that are available within the current CLI context, enter a question mark (`?`) at the command prompt.
- To display a list of the available commands that start with a particular character set, enter the abbreviated command immediately followed by a question mark (`?`).
- To display help for a command's legal arguments, enter a question mark (`?`) in place of an argument at the command prompt.



Note The question mark (`?`) is not echoed back to the console.

Syntax:

```
?
abbreviated_command ?
command [arguments] ?
```

Example: `> ?`

Cloud-Delivered Firewall Management Center CLI show commands

Show commands provide information about the state of the appliance. These commands do not change the operational mode of the appliance and running them has minimal impact on system operation.

version command

Displays the product version and build, the UUID, and other information.

Syntax:

```
show version
```

Example:

```
> show version
-------------------[ fmc-austin ]--------------------
Model                     : Cisco Secure Firewall Management Center for VMware (66) Version
```

```
 7.6.0 (Build 1385)
UUID                        : a904b8b2-ca9a-11ee-a583-5e804c16b2fd
Rules update version        : 2024-05-13-001-vrt
LSP version                 : lsp-rel-20240513-1955
VDB version                 : 380
----------------------------------------------------
```

Cloud-Delivered Firewall Management Center CLI configuration commands

The configuration commands enable the user to configure and manage the system. These commands affect system operation.

password command

Allows the current CLI user to change their password.

Caution For system security reasons, we strongly recommend that you do not establish Linux shell users in addition to the pre-defined **admin** on any appliance.

After issuing the command, the CLI prompts the user for their current (or old) password, then prompts the user to enter the new password twice.

Syntax:

```
configure password
```

Example:

```
> configure password
Changing password for admin.
(current) UNIX password:
New UNIX password:
Retype new UNIX password:
passwd: password updated successfully
```

Cloud-Delivered Firewall Management Center CLI system commands

The system commands enable the user to manage system-wide files and access control settings.

Following sections describes the Cloud-Delivered Firewall Management Center CLI system commands:

generate-troubleshoot command

Generates troubleshooting data for analysis by Cisco.

Syntax:

```
system generate-troubleshoot option1 optionN
```

Where options are one or more of the following, space-separated:

- `SYS`: System Configuration, Policy, and Logs
- `PER`: Hardware Performance and Logs
- `NET`: Interface and Network Related Data
- `UPG`: Upgrade Data and Logs
- `SNT`: Snort Performance and Configuration
- `DES`: Detection Configuration, Policy, and Logs
- `VDB`: Discover, Awareness, VDB Data, and Logs
- `DBO`: All Database Data
- `LOG`: All Log Data
- `NMP`: Network Map Information
- `ALL`: Run all of the following options.

Example:

```
> system generate-troubleshoot VDB NMP
starting /usr/local/sf/bin/sf_troubleshoot.pl…
Please, be patient. This may take several minutes.
The troubleshoot options codes specified are VDB,NMP.
Getting filenames from [usr/local/sf/etc/db_updates/index]
Getting filenames from [usr/local/sf/etc/db_updates/base-6.2.3]
Troubleshooting information successfully created at
/var/common/results-06-14-2018—222027.tar.gz
```

lockdown command

Removes the `expert` command and access to the Linux shell on the device.



Caution This command is irreversible without a hotfix from Support. Use with care.

Syntax: `system lockdown`

Example:

```
> system lockdown
```

reboot command

Reboots the appliance.

Syntax: `system reboot`

Example:

```
> system reboot
```

restart command

Restarts the appliance application.

Syntax: `system restart`

Example:

```
> system restart
```

show max-logins command

Displays the maximum number of parallel logins for one user.

Syntax: `show max-logins`

Command default: The default value for the number of parallel/concurrent sessions for a single user is zero.

Example:

```
> show max-logins
```

configure max-logins command

Configures the maximum number of parallel logins for one user.

Syntax: `configure max-logins`

Example:

```
> configure max-logins [0-1024]
```

shutdown command

Shuts down the appliance.

Syntax: `system shutdown`

Example:

```
> system shutdown
```



CHAPTER 107

Communication Ports and Internet Access

- Communication ports for managed devices, on page 3115
- Internet access requirements for managed devices, on page 3116
- Internet resources accessed by your browser, on page 3118

Communication ports for managed devices

Managed devices use these ports to communicate. For deployments behind a network barrier—like an edge firewall—make sure you allow traffic on the required ports. Note that ports not required for essential or default operations remain closed until needed by a configuration or feature.

Inbound ports for managed devices

Managed devices accept inbound traffic on these ports.

Table 165: Inbound ports for managed devices

Inbound port	Protocol/Feature	Details
Required for specific configurations or features		
22/tcp	SSH	Secure remote connections to the appliance.
161/udp	SNMP	Allow access to MIBs via SNMP polling.
443/tcp	Remote access VPN (SSL)	Allow secure VPN connections to your network from remote users.
443/udp	Remote access VPN (DTLS)	Allow secure VPN connections to your network from remote users.
500/udp 4500/udp	Remote access VPN (IKEv2) and site-to-site VPN	Allow secure VPN connections to your network from remote users and remote VPN peers.
885/tcp	Captive portal	Communicate with a captive portal identity source.

Outbound ports for managed devices

Managed devices initiate outbound traffic on these ports. Managed devices also use ephemeral source ports for TCP and UDP traffic that they initiate. Make sure intervening network barriers allow response traffic for these connections.

Table 166: Outbound ports for managed devices

Outbound port	Protocol/Feature	Details
Required for initial setup		
53/tcp 53/udp	DNS	DNS
123/udp	NTP	Synchronize time.
443/tcp	HTTPS	Send and receive data from the internet; see Internet access requirements for managed devices, on page 3116 for a list of resources that the device needs to access. Also accepts connections on this port.
8305/tcp	Appliance communications	Securely communicate with the Cloud-Delivered Firewall Management Center.
Required for specific configurations or features		
67/udp 68/udp	DHCP	DHCP
162/udp	SNMP	Send SNMP alerts to a remote trap server.
1812/udp 1813/udp	RADIUS	Communicate with a RADIUS server for external authentication and accounting. Configurable.
389/tcp 636/tcp	LDAP	Communicate with an LDAP server for external authentication. Configurable.
514/udp	Syslog (audit logging)	Send audit logs to a remote syslog server, when TLS is not configured.
8514/udp	Secure Network Analytics Manager	Send syslog messages to Secure Network Analytics using Security Analytics and Logging (On Premises).
8989/tcp	Cisco Support Diagnostics	Port 8989/tcp is used only for local/on-box communication between Cisco Support Diagnostics and SSEConnector. Cisco Support Diagnostics cloud/CDX communication is proxied through SSEConnector and uses HTTPS/WSS on 443/tcp externally.

Internet access requirements for managed devices

Internet access to the listed resources supports essential operations and configured features in your deployment.

For deployments behind a network barrier—like an edge firewall—make sure you allow traffic on the required ports to the listed resources.

Internet resources accessed by managed devices

This table identifies the internet resources that managed devices access.

Table 167: Internet resources for managed devices

Feature	Reason	HA/Clustering	Port	Resource
Required for initial setup				
DNS	DNS	All units communicate with the DNS server.	53/tcp 53/udp	Default OpenDNS servers: 208.67.222.222 and 208.67.220.220 2620:119:35::35
NTP	Synchronize time. Not supported with a proxy server.	All units communicate with the NTP server.	123/udp	Default NTP servers: 0.sourcefire.pool.ntp.org 1.sourcefire.pool.ntp.org 2.sourcefire.pool.ntp.org 3.sourcefire.pool.ntp.org
Required for general operations				
CA certificate bundles	Managed devices query for new CA certificates daily at a system-defined time. The local CA bundle contains certificates for accessing several Cisco services.	Each unit downloads its own certificates.	443/tcp	cisco.com/security/pki
Cisco Support Diagnostics	Cisco Support Diagnostics receives authorized CSD/CDX troubleshooting requests and uploads generated troubleshooting files through SSEConnector. External Cisco cloud connectivity utilizes HTTPS/WSS on 443/tcp.	All units communicate.	443/tcp	api-sse.cisco.com
Required for specific configurations or features				
Malware Defense	Submit files for dynamic analysis.	All units submit files.	443/tcp	fmc.api.threatgrid.com fmc.api.threatgrid.eu

Feature	Reason	HA/Clustering	Port	Resource
Upgrades	Download upgrades directly to managed devices. Tests the connection once a week.	Upgrade packages do not sync. Each unit must get its own.	443/tcp	cdo-fdm-images3-us-west-2.amazonaws.com
Umbrella DNS	Direct DNS queries to Cisco Umbrella for validation and policy enforcement.	All units communicate.	443/tcp	api.opendns.com

Security Cloud Control regional IP addresses

To onboard and manage devices using zero-touch provisioning, allow inbound access on port 443 (or whichever port you configured for device management) from the Security Cloud Control IP addresses in your region.

Table 168: Security Cloud Control IP addresses by region

Security Cloud Control region	IP addresses
Asia-Pacific-Japan (APJ)	54.199.195.111 52.199.243.0
Australia (AUS)	13.55.73.159 13.238.226.118
Europe, the Middle East, or Africa (EMEA)	35.157.12.126 35.157.12.15
India (IN)	35.154.115.175 13.201.213.99
United States (US)	52.34.234.2 52.36.70.147

Internet resources accessed by your browser

When you use the Cloud-Delivered Firewall Management Center, your browser may contact Google (google.com) or Amplitude (amplitude.com) web analytics servers to send non-personally-identifiable usage data to Cisco. For browser compatibility information, refer to the *Cisco Secure Firewall Management Center Compatibility Guide*.